arXiv:1704.02778v1 [stat.ME] 10 Apr 2017

Multiscale Bayesian State Space Model for Granger Causality Analysis of Brain Signal

Sezen Cekic
Methodology and Data Analysis, Department of Psychology,
University of Geneva,
Didier Grandjean
Neuroscience of Emotion and Affective Dynamics Lab,
Department of Psychology,
University of Geneva,
and
Olivier Renaud
Methodology and Data Analysis, Department of Psychology,
University of Geneva

April 11, 2017

Abstract

This article concerns the modelling of time-varying and frequency-specific relationships between two signals, with a focus on intracerebral signals measuring neural activities. Many researchers in neuroscience would like to assess what is called a frequency Granger causality that may vary in time to evalute the functionnal connections between two brain regions during a task. We propose the use of an adaptive Kalman filter type of estimator of a linear Gaussian vector autoregressive model with coefficients evolving over time. The estimation procedure is achieved through variational Bayesian approximation and can be extended for multiple trials. This Bayesian State Space (BSS) model provides a dynamical Granger-causality statistic

that is quite natural. The Bayesian nature of the model provides a criterion for model order selection and allows us to include prior knowledge in the model. We propose to extend the BSS model to include the *à trous* Haar decomposition. This wavelet-based forecasting method, based on a multiple resolution decomposition of the signal using the redundant *à trous* wavelet transform, captures short- and long-range dependencies between signals and is further used to derive the desired dynamical and frequency-specific Granger-causality statistic. The application of these models to local field potential data recorded during a psychological experimental task shows the complex frequency based cross-talk between amygdala and medial orbito-frontal cortex.

Keywords: Time-frequency, variational methods, *à trous* Haar wavelets, multiple trials, nonstationarity, neuroscience data.

1 Introduction

In this article we present a new estimation and prediction methodology for multivariate nonstationary time series, in order to get a Granger-type dynamic causal statistic that allows us to model nonstationary brain dynamics between two series of data recorded in specific brain areas.

In many neuroscientific experiments, data are recorded in an experimental situation where stimuli are presented at fixed times and are expected to induce a reaction. For psychologists and neuroscientists, being able to model and explain the dynamics of the functional and effective links between neural and behavioural signals recorded during the experiment is of primary interest. They often have strong prior hypotheses about these causal links and therefore need reliable statistical tools to draw valid conclusions.

We propose to derive a dynamical Granger-causal statistic from a linear Gaussian vector autoregressive (VAR) model with coefficients evolving over time according to a linear dynamical system. Given that this model is strongly over-parametrized, we propose to place it in a Bayesian framework and to use the variational method to estimate all the densities. The variational Bayesian methodology (Beal, 2003) estimates of all the necessary quantities and moreover allows us to deal with datasets containing several trials (or epochs). The Bayesian nature of the model offers a natural criterion for model order selection, which is a very important issue in time series modelling. In Section 3, we describe our Bayesian state space (BSS) model and discuss its technical specificities and the estimation procedure. In Section 5, we propose an extension of the BSS model, called the multiscale Bayesian state space (MSBSS) model, which is based on the *à trous* multiscale wavelet transform. The latter approach has never been used in this context of time-varying VAR coefficient estimate and we will show that it allows better estimate of short- and long-range specific

dependencies between recorded signals, and moreover that it offers a very simple way to deal with time-frequency uncertainty bounds. In Section 6, we will present a Bayesian dynamical Granger-causality statistic based on the time-varying estimated VAR coefficients and in Section 7, we present simulation studies to assess our proposed methodology.

The major contribution of this article is to provide a statistical methodology that can be used to address dynamic Granger causality for neuroscience data. We offer a unified framework, based from several previous works from different research areas including engineering, statistics and methodological neuroscience, and test it on simulated and real data.

2 Granger causality

The question of how to operationally formalize and test for causality is a deep fundamental and philosophical problem. A mathematical solution, which relies on the causal nature of predictability, was provided in the 60's by the economist Clive Granger and was latter coined "Granger causality". According to Granger (1969), if a signal X "Granger-causes" a signal Y, then the history of X should contains information that helps to predict Y above and beyond the information contained in the history of Y alone. The axiomatic imposition of a temporal ordering is the crucial element that enables us to interpret such dependence as causal.

As its causal nature relies on prediction, Granger causality does not necessarily means "true causality", for instance, if the two studied processes indeed are commonly driven by a third one, one might reject the null hypothesis of non-Granger causality between the two signals. Of course, Granger causality rely on a suitable preliminary variables selection procedure as well as a good model estimation of the underlying VAR model as described

in the next section.

Application of the original Granger causality relies on linear Gaussian autoregressive models and therefore on the assumption that the analysed system is covariance stationary (Granger, 1969). In this context, a common way to test for Granger causality between two series is to estimate a vector autoregressive model (VAR) and then test the significance of the off diagonal coefficients of interest (Granger, 1969; Lütkepohl, 2005).

However, in neuroscience data are usually nonstationary and this characteristic is moreover of interest. We therefore want to derive a causality-statistic that allows us to capture the causal structure differentially for each time. This will be achieved by extending Basic causality statistic letting the VAR model evolve in time. This time-varying VAR model is strongly over-parametrized. In addition, we need: an objective criterion for model order selection, to deal with datasets containing several trials (or epochs) and to capture short- and long-range causal dependencies between signals due to specific frequency characteristics.

3 The Bayesian State Space Model

We propose to write the dynamic VAR model in a state space form which allows us to rewrite the dynamic VAR model as an observation equation in which the dynamic VAR coefficients are driven by the state equation (this modelling proposal was previously made by Cassidy, 2002). This leads to the following system of equations:

$$\begin{cases} \varphi_{t+1} = A\varphi_t + w_t & w_t \sim \mathcal{N}_k(0, Q) \\ Z_t = C_t\varphi_t + v_t & v_t \sim \mathcal{N}_d(0, R) \end{cases} \quad \text{with} \quad \begin{cases} \varphi_t & = \text{vec}[\vartheta_{1(t)}, \vartheta_{2(t)}, .., \vartheta_{p(t)}]', \\ Z_t & = (Y_t \; X_t)', \\ C_t\varphi_t & = \sum_{j=1}^{p} \vartheta_{j(t)}(Y_{t-j} \; X_{t-j})', \end{cases} \tag{1}$$

where $Z_t = (Y_t\ X_t)'$ is the value of the $d = 2$ signals at time t, $\vartheta_{j(t)}$ are the time-varying VAR coefficients (up to order p) and the vector φ_t of size $k = pd^2$ contains all the time-varying VAR coefficients that have to be estimated for the time t. The matrix A is the transition matrix of the state vector φ_t, Q is the $k \times k$ variance-covariance matrix of the state equation, and R is the $d \times d$ variance-covariance matrix of the observation equation. We will use the notation φ_1^T and Z_1^T to denote the entire set of values from $t = 1$ to $t = T$.

With a slight abuse of notation, we can write $p(Z_t|C_t, \varphi_t, R) = \mathcal{N}_d(C_t\varphi_t, R)$, $p(\varphi_t|A, \varphi_{t-1}, Q) = \mathcal{N}_k(A\varphi_{t-1}, Q)$ and $p(\varphi_1) = \mathcal{N}_k(\mu_1, \Sigma_1)$, which are respectively the observation, the state and the initial state densities.

Given the Bayesian framework, we are interested in obtaining the posterior distribution of the unknowns of our model. In simple cases, these distributions can be obtained analytically but here, the curse of dimensionality makes the computation intractable and forces us to rely on approximation techniques.

3.1 Elements of variational Bayes

Variational approximation is not widely known within the statistical community dominated by Monte Carlo approximation methods as well as Laplace approximation methods. Variational approximation however is much faster than Monte Carlo, especially for large models, and allows us to deal with models containing a very large amount of parameters.

For sake of clarity, we will define the set of unknown parameters as Ω_1^b, so the full set of unknowns for model (1) is $\{\varphi_1^T; \Omega_1^b\}$.

Our target quantity is the posterior distribution which is an intractable integral of very high dimension and the variational approach allows us to approximate this posterior density $p(\varphi_1^T, \Omega_1^b|Z_1^T)$ by a variational posterior density $q(\varphi_1^T, \Omega_1^b|Z_1^T)$ that will be selected to be optimal according to the Kullback–Leibler distance dissimilarity criterion (Beal, 2003).

3.1.1 Learning rules

By simple algebra, we will decompose the evidence of the model $p(Z_1^T)$ by inserting the variational density q. We have

$$\log p(Z_1^T) = \mathrm{KL}\big(q(\varphi_1^T, \Omega_1^b | Z_1^T) \| p(\varphi_1^T, \Omega_1^b | Z_1^T)\big) - \left\langle \log \frac{q(\varphi_1^T, \Omega_1^b | Z_1^T)}{p(\varphi_1^T, \Omega_1^b, Z_1^T)} \right\rangle_{q(\varphi_1^T, \Omega_1^b | Z_1^T)}, \qquad (2)$$

where $\langle . \rangle$ denotes expectation and its subscript denotes the density used for this expectation.

By the positiveness of the KL distance, we have obtained a lower bound for the logarithm of the evidence

$$\log p(Z_1^T) \geq - \left\langle \log \frac{q(\varphi_1^T, \Omega_1^b | Z_1^T)}{p(\varphi_1^T, \Omega_1^b, Z_1^T)} \right\rangle_{q(\varphi_1^T, \Omega_1^b | Z_1^T)} := \mathrm{F}, \qquad (3)$$

where F is called the negative free energy. The aim is thus to find a density $q(\varphi_1^T, \Omega_1^b | Z_1^T)$ for which the integrals in F are tractable and which is close to $p(\varphi_1^T, \Omega_1^b | Z_1^T)$.

3.1.2 Mean-field approximation

The choice underlying the variational Bayesian methodology is to allow the approximating density $q(\varphi_1^T, \Omega_1^b | Z_1^T)$ to factorize over groups of parameters. We will suppose here that the approximating density factorizes as

$$q(\varphi_1^T, \Omega_1^b | Z_1^T) = q(\varphi_1^T | Z_1^T) \prod_{j=1}^{b} q(\Omega_j | Z_1^T), \qquad (4)$$

and the lower bound for the model evidence can be rewritten based on this factorization.

3.1.3 Variational EM algorithm

With the use of the calculus of variations, it can be shown that under assumption (4), the variational distributions $q^*(\varphi_1^T | Z_1^T)$ and $q^*(\Omega_j | Z_1^T)$, that maximize the functional F, can be

expressed and therefore maximized in an iterative way (Beal, 2003). First,

$$q^*(\varphi_1^T|Z_1^T)^{(l+1)} \propto \exp\left\langle \log p(\varphi_1^T|\Omega_1^k, Z_1^T)\right\rangle_{q(\Omega_1^b|Z_1^T)}^{(l)}, \tag{5}$$

where superscript (l) denotes the iteration number. The other steps for $m = 1, \ldots, b$ are

$$q^*(\Omega_m|Z_1^T)^{(l+1)} \propto \exp\left\langle \log p(\Omega_1^k|\varphi_1^T, Z_1^T)\right\rangle_{-\Omega_m}^{(l)}, \tag{6}$$

where $\langle . \rangle_{-\Omega_m}^{(l)}$ is the expectation over all the distributions at iteration l except $q(\Omega_m|Z_1^T)^{(l)}$. See Beal (2003), and Ostwald et al. (2014) for all proofs.

Thus, if we set $q(\varphi_1^T|Z_1^T)$ equal to $q^*(\varphi_1^T|Z_1^T)$ and $q(\Omega_m|Z_1^T)$ equal to $q^*(\Omega_m|Z_1^T)$, the lower bound for the model evidence F is maximized under the (4) constraint. The form of equations (5) and (6) defines a circular dependence which explains the use of an iterative algorithm whose convergence can be assessed by monitoring the relative increase of F. This iterative algorithm is known as the variational Bayesian expectation-maximisation algorithm (Beal, 2003).

3.1.4 Conjugate exponential

The optimal form for $q^*(\varphi_1^T|Z_1^T)$ and $q^*(\Omega_1^b|Z_1^T)$ of course depends on the choice of the prior distributions $p(\varphi_1^T)$ and $p(\Omega_m)$. The analytical form of $q^*(.|Z_1^T)$ will be of the same distributional form as the prior distributions $p(.)$, if the complete-data likelihood $p(\varphi_1^T, Z_1^T|\Omega_1^b)$ is part of the exponential family and if the hidden and parameter prior distributions $p(\varphi_1^T)$ and $p(\Omega_1^b)$ are conjugate to this complete-data likelihood (Beal, 2003). As stated in Beal (2003, p. 160), with suitable priors, the state space model with unknown parameters is in the conjugate-exponential family.

3.2 The variational evidence lower bound

In Bayesian analysis, the evidence provides a natural criterion for model selection. As observed in equation (3), the variational Bayesian methodology has the advantage of providing a quantity F that is a lower bound for the evidence of the model and that can be computed efficiently. This leads to a natural criterion for model order selection which is crucial to estimate our strongly over-parametrized time-varying VAR model. We can therefore perform model selection by comparing the F quantities computed for each model order m_p and select the model that exhibits the highest F_{m_p}. The specific analytic form of F_{m_p} for our model is derived in Appendix A.

4 Model specification

We will now describe the fully hierarchical Bayesian model that we propose. The related complete DAG can be found in Figure 1. Recalling the model (1), the complete data likelihood $p(\varphi_1^T, Z_1^T|\Omega_1^b)$ is Gaussian and thus falls into the class of exponential family distributions. We will now describe the model by defining conjugate priors over the model parameters. As stated in Section 3.1.4, the corresponding variational approximate posterior distributions will then be of the same analytical form.

4.1 Prior distributions

4.1.1 Prior for φ_1^T

The state space structure of the model (1) illustrated in Figure 1 leads to the conditional prior distributions for the hidden state φ_1^T defined in Section 3 and the prior mean μ_1 and variance Σ_1 for φ_1 are set to $1_k \times 0$ and $I_k \times 0.1$ respectively.

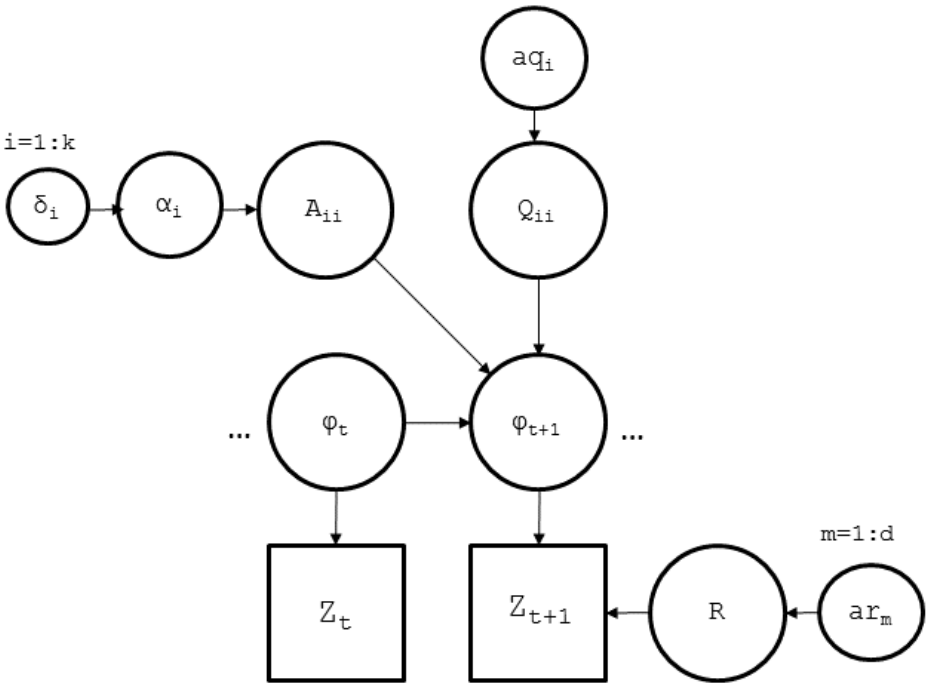

Figure 1: DAG representation of the full model given in Section 4.1.

4.1.2 Prior for $\Omega_1 = \{A, \alpha, \delta\}$

We propose a diagonal structure for the A matrix, meaning that a specific causal coefficient at time t will only depend on its own past value plus a white noise and not on the past values of other VAR coefficients. This assumption seems reasonable in terms of brain connectivity, where the dynamic of a particular causal input can be assumed to be driven by its own trajectory and not by that of other causal or auto-causal inputs. To satisfy the conjugacy condition of stated in Section 3.1.4, we place a Gaussian prior on each diagonal entry of the A matrix:

$$p(A|\alpha) = \prod_{i=1}^{k} \mathcal{N}_1(A_{ii}|m_{A_{ii}}, \alpha_i), \tag{7}$$

where α_i is the hyperparameter variance of the specific diagonal element A_{ii} whose distribution will be discussed below. We choose to impose a conservative prior for A by setting each location hyperparameter $\{m_{A_{ii}}\}_{i=1}^{k}$ to 0.9.

As we have no subjective input for the variances α_i, it is desirable to have priors that exhibit very little information. The non-informativeness of a variance prior is a sensitive topic in Bayesian estimation procedures.

The main recommendation in Gelman (2006) is to use half-t priors on standard deviation parameters to achieve arbitrarily high noninformativeness. The definition of the half-t distribution can be found e.g. in Wand et al. (2011).

Half-t priors may appear to be computationally harder to work with than classical inverse-gamma priors, but Wand et al. (2011) showed that the half-t distribution can be written as a scaled mixture of inverse-gamma distributions. So the half-t prior is obtained for each $\alpha_i^{0.5}$ element through the auxiliary variable construction

$$p(\alpha_i|\delta_i) = \mathcal{IG}(c_{p_i}, b_{p_i}) = \mathcal{IG}(\frac{1}{2}, \frac{1}{\delta_i}), \quad p(\delta_i) = \mathcal{IG}(\kappa_{p_i}, \beta_{p_i}) = \mathcal{IG}(\frac{1}{2}, \frac{1}{D_i^2}), \quad i = 1, \ldots, k, \tag{8}$$

which ensures that $p(\alpha_i^{0.5}) = \text{Half-t}(1, D_i)$. The hierarchical representation in equation (8) respects the conditional conjugacy condition due to the conjugacy properties of the inverse-gamma distribution. All the theoretical justifications can be found in Wand et al. (2011).

By its noninformativeness property, this half-t prior choice will let each variance element $\{\alpha_i\}_{i=1}^k$ get an appropriately high or low value, thereby allowing them to play the role of shrinkage parameters for the distribution of A_{ii}. In our specific case, it will allow us to be conservative and to tend to avoid an erroneous causality assessment.

4.1.3 Prior for $\Omega_2 = \{Q, a_q\}$

The accuracy of the Markovian conditional distribution followed by the VAR coefficients φ_t is established through the variance-covariance matrix Q. Throughout this work we assume that this matrix is diagonal (for the same theoretical reason discussed for the A matrix) and that all these elements are equal. In fact there is no reason to assume that the variability of one VAR coefficient is different from that of another. This choice is also motivated by the concern to limit the number of parameters to estimate.

In order to satisfy the conjugacy conditions and for the same theoretical reason mentioned earlier, we set the same weakly-informative half-t prior distribution for the single standard deviation parameter $Q_{ii}^{0.5}$ of the diagonal Q matrix:

$$p(Q_{ii}|a_q) = \mathcal{IG}(n_p, d_p) = \mathcal{IG}(\frac{1}{2}, \frac{1}{a_q}), \quad p(a_q) = \mathcal{IG}(a_{qp}, b_{qp}) = \mathcal{IG}(\frac{1}{2}, \frac{1}{A_q^2}), \tag{9}$$

which again ensures that $p(Q_{ii}^{0.5}) = \text{Half-t}(1, A_q)$.

4.1.4 Prior for $\Omega_3 = \{R, a_r\}$

The variance-covariance matrix R of the observed equation is not supposed to be diagonal due to the theoretical interdependence between brain signals modelled in the system.

We impose for R a generalisation of the multivariate case of the Half-t prior used for α and Q_{ii}. Huang et al. (2013) derived this prior and explained that with a suitable hyperparameter choice, it induces an Half-t distribution for each standard deviation term corresponding to the diagonal of R, as well as a marginal uniform distribution for all correlations. Explicitly we have

$$\begin{aligned} p(R|a_{r_1}, \ldots, a_{r_d}) = \mathcal{IW}_d\big(r_p, B_p\big) = \mathcal{IW}_d\big(\nu + d - 1, 2\,\nu\,\mathrm{diag}[\frac{1}{a_{r_1}}, \ldots, \frac{1}{a_{r_d}}]\big), \\ p(a_{r_i}) = \mathcal{IG}(a_{pr}, b_{pr}) = \mathcal{IG}(\frac{1}{2}, \frac{1}{A_R^2}), \qquad i = 1, \ldots, d, \end{aligned} \tag{10}$$

where diag() denotes a diagonal matrix. This hierarchical structure ensures that each diagonal element $R_{i,i}$ is distributed as $p(R_{i,i}^{0.5}) = \text{Half-t}(\nu, A_R)$, and the particular choice of $\nu = 2$ leads to marginal uniform distributions over $[-1; 1]$ for all correlation terms (Huang et al., 2013).

Lastly, we note that the selected prior distributions impose the choice of hyperparameters $D_1, \ldots, D_k$, A_q, ν and A_R. We will choose $\nu = 2$ for the uniform property to hold. For the remaining hyperparameters, the larger they are the less informative the priors are. We follow Menictas and Wand (2013) and set them to 10^5.

4.2 Update equations

We now provide details on the variational Bayesian iterative algorithm based on the results presented in Section 3.1.3 and factorization chosen in Section 3.1.2.

4.2.1 Variational E-step

The optimal variational posterior distribution for the hidden state sequence φ_1^T is multivariate Gaussian at each time t:

$$q^*(\varphi_t|Z_1^T) = \mathcal{N}_k(\varphi_t|\mu_t, \Sigma_t), \tag{11}$$

where the sufficient statistics $\{\mu_t; \Sigma_t\}_{t=1}^T$, as well as the cross-moments $\{\mu_t\mu_{t-1}; \Sigma_{t,t-1}\}_{t=2}^T$, are obtained by the well-known Kalman–Rauch–Tung–Striebel (KRTS) smoother algorithm (Kalman, 1960; Kalman and Bucy, 1961). KRTS smoother recursive equations applied to the suitable augmented system of equations is discuss and derived in Appendix D.

4.2.2 Variational M-step

All the derivations can be found in Appendices E, F, G, H, I, J and K. Here we give the resulting distributions.

Update for $\Omega_1 = \{A, \alpha, \delta\}$ Each hyper-hyperparameter $\{\delta_i\}_{1=1}^k$ has an inverse-gamma variational posterior distribution $q^*(\delta_i|\kappa_i, \beta_i)^{(l+1)}$ with parameters $\kappa_i = \kappa_{p_i} + c_{p_i}, \quad \beta_i = \beta_{p_i} + \left\langle \alpha_i^{-1} \right\rangle$, where $\langle . \rangle$ denotes the expectation with respect to the variational posterior distribution of the corresponding variable at iteration l. The complete derivation can be found in Appendix E.

The variance α_i of each diagonal entry A_{ii} of the A matrix has an inverse-gamma variational posterior distribution by conjugacy $q^*(\alpha_i|c_i, b_i)^{(l+1)}$ with parameters $c_i = c_{p_i} + 1/2, \quad b_i = \left\langle \delta_i^{-1} \right\rangle + (\theta_i + (\psi_i\theta_i)^2 + m_{A_{ii}}^2 - 2m_{A_{ii}}\psi_i\theta_i)/2$, where $\langle \delta_i \rangle$ denotes the expectation of δ_i under the posterior distribution $q^*(\delta|Z_1^T)^{(l)}$ defined in Appendix E. The complete derivation can be found in Appendix F.

The prior densities for the diagonal elements of the matrix A being k independent Gaussian variables A_{ii}, as described in equation (7), the variational posterior $q^*(A_{ii}|\mu_{A_{ii}}, \sigma^2_{A_{ii}})^{(l+1)}$ is therefore univariate Gaussian by conjugacy, with mean and variance $\mu_{A_{ii}} = \left[\left\langle \alpha_i^{-1} \right\rangle m_{A_{ii}} + M_{3\{i,i\}} \langle Q_{ii} \rangle \right] \left[\langle Q_{ii} \rangle M_{2\{i,i\}} + \left\langle \alpha_i^{-1} \right\rangle \right]^{-1}, \sigma^2_{A_{ii}} = \left[\langle Q_{ii} \rangle M_{2\{i,i\}} + \left\langle \alpha_i^{-1} \right\rangle \right]^{-1}$, where M_2, M_3 are defined in Appendix D. The derivation of $q^*(A_{ii}|\mu_{A_{ii}}, \sigma^2_{A_{ii}})^{(l+1)}$ can be found in Appendix G.

Update for $\Omega_2 = \{Q, a_q\}$ Recalling that the hyperparameter a_q *a priori* follows an inverse-gamma distribution, as described in equation (9), the variational posterior distribution $q^*(a_q|a_{qq}, b_{qq})^{(l+1)}$ is accordingly also inverse-gamma distributed but with shape and scale parameters $a_{qq} = a_{qp} + n_p, b_{qq} = b_{qp} + \langle q^{-1} \rangle$. The complete derivation can be found in Appendix H.

The matrix Q is proportional to the identity, therefore with a single random element Q_{ii} that follows *a priori* an inverse-gamma distribution as described in equation (9). The variational posterior distribution $q^*(Q_{ii}|n, d)^{(l+1)}$ is accordingly also inverse-gamma but with shape and scale parameters $n = n_p + k(T-1)/2, d = \left\langle a_q^{-1} \right\rangle + \Gamma/2$. The entire derivation for equation $q^*(Q_{ii}|n, d)^{(l+1)}$ can be found in Appendix I.

Update for $\Omega_3 = \{R, a_r\}$ Recalling that each hyperparameter $\{a_{r_i}\}_{i=1}^d$ *a priori* follows an inverse-gamma distribution as described in equation (10), the variational posterior distribution $q^*(a_{r_i}|a_{qr}, b_{qr})^{(l+1)}$ is accordingly an inverse-gamma with shape and scale parameters $a_{qr} = a_{pr} + r_p/2, \quad b_{qr} = b_{pr} + \nu \left\langle R^{-1}_{\{i,i\}} \right\rangle, \quad$ where $\quad \langle R^{-1} \rangle = r \langle B^{-1} \rangle$. The derivation for $q^*(a_{r_i}|a_{qr}, b_{qr})^{(l+1)}$ can be found in Appendix J. The variational posterior distribution for the variance-covariance matrix R is inverse-Wishart by conjugacy $\mathcal{IW}_d(R|r, B)$, with corresponding parameters $B = 2\,\nu\,\mathrm{diag}[\frac{1}{\langle a_{r_1} \rangle}, \ldots, \frac{1}{\langle a_{r_d} \rangle}] + \sum_{t=1}^{T} [(Z_t -$

$C_t\mu_t)(Z_t - C_t\mu_t) + C_t\Sigma_t C_t^{'} R], r = r_p + T$. The complete derivation can be found in Appendix K.

4.3 Multiple trials

One important contribution of the present article is to show that our model can be modified to deal with N conditionally independent sequences $\{Z_{1\ (j)}^T\}_{j=1}^N$ which are supposed to have the same hidden state. This reflects the case that arises during an event-related experimental paradigm, where many trials on the same condition are measured.

In Beal (2003) and Cassidy (2002), this extension is solved by first estimating the necessary sufficient statistics for each sequence independently in the E-step, and then by averaging these statistics to get only one set of sufficient statistics, which is representative of the entire set of independent sequences before performing the M-step. However this approach does not take into account the complex dependence of the variational posterior $q(\varphi_1^T|\{Z_{1\ (j)}^T\}_{j=1}^N)$ on the whole dataset $\{Z_{1\ (j)}^T\}_{j=1}^N$. In fact, all the computations done so far can be adapted to multiple trials. A detailed derivation of the model and related variational posteriors distributions in a multiple trials setting can be founded in Appendix L.

5 The Multiscale Bayesian State Space Model

In a neuroscience context, an important limitation of many models, including the BSS model, is the inability to capture both the short- and long-range possible causal dependencies between signals. The result should be as invariant as possible to arbitrary choices like the chosen sampling frequency or added time-lag in the prediction model.

Based on this consideration, and because Granger causality is based on predictive ability, if the auto-causal information contained in the history of the predicted signal Y_t in model

(1) is not well represented in C_t, the Granger-causality evaluation, which is based on the predictability improvement of Y_t by adding the information contained in the history of the second variable X, may be spuriously assessed as significant. On the other hand, if the predictive ability of the history of the causal signal X is not informative enough, we can miss some crucial information about causal interdependencies between signals Y and X.

In this section, we propose a new solution that has the ability to appropriately select the short- and long-term causal histories of Y and X by combining the BSS model with the *à trous* multiscale decomposition methodology (Renaud et al., 2003).

5.1 The *à trous* Haar wavelets transform

For that purpose, we will perform the *à trous* multiscale decomposition of signals Y_t and X_t contained in Z_t in model (1), in order to use these quantities as predictive histories in the matrices $\{C_t\}_{t=1}^{T}$. The reader is referred to Renaud et al. (2003) and references therein for a complete overview of the method. We define $w_{j,t}^{y}$ as the *à trous* wavelet coefficient of the signal Y at time t for scale j. Since we will use the *à trous* wavelet coefficients for prediction, they should not be based on future values and the only family that satisfies this constraint is the *à trous* Haar wavelet transform. Then $S_{j+1}(t) = [S_j(t) + S_j(t - 2^j)]/2, \quad w_{j+1}(t) = S_j(t) - S_{j+1}(t), \quad t = 1, \ldots, T, \quad j = 1, \ldots, J$, where the finest scale is the original series $S_0(t) = Y_t$.

5.2 The multiscale Bayesian state space model

Based on the derivation in Section 5.1, we can modify the VAR model in equation (1). We keep the quantity to predict Y_t in the time domain, but the histories of the series Y and X will equal the *à trous* Haar wavelet transforms of series Y_t and X_t respectively. We can

therefore define the set C_t^w that contains the decompositions of the histories of series Y_t and X_t as

$$\begin{aligned} C_t^w = &\{w^y_{j,t-1-2^j(k-1)}\}_{j=1,\dots,J,k=1,\dots,p_j}, \{s^y_{J,t-1-2^J(k-1)}\}_{k=1,\dots,p_{J+1}}, \\ &\{w^x_{j,t-1-2^j(k-1)}\}_{j=1,\dots,J,k=1,\dots,p_j}, \{s^x_{J,t-1-2^J(k-1)}\}_{k=1,\dots,p_{J+1}}, \end{aligned} \tag{12}$$

and therefore adapt equation (1) replacing C_t with C_t^w.

In particular, we underline that for any J, for a suitable choice of $p_1, \dots, p_{J+1}$, C_t^w is an orthogonal transform of C_t and therefore in this specific case BSS and MSBSS models are equivalent. For any choice of the p's, all results in Section 3 and all the estimation procedures in Section 4 can be applied. The only difference is that the matrix C_t is replaced by C_t^w, and that the dimension of φ_t is changed accordingly.

The *à trous* extension can thus be viewed as a generalisation of the BSS model that contains information relative to the frequencies (defined by the wavelet scales). The model order p_j for the different scales as well as the number of scales J to be taken in the model must now be selected. As argued in Renaud et al. (2003), the relative non-overlapping frequencies used in each scale motivate an independent selection of the model order p_j for each scale. Defining a maximal scale decomposition $J_{\max}$ and a maximum model order for each scale $p_{\max}$, the model order selection procedure set in Section 3.2 is iteratively applied to select p_j in a stepwise manner for each scale from 1 to $J_{\max}$. The free-energy quantities related to the $J_{\max}$ models are then compared and the J that exhibits the highest free energy is selected. The model order for the smooth p_{J+1} is finally selected.

This multiresolution methodology presents several significant benefits. As the wavelet coefficients w_j are directly related to a specific frequency band, this enables us to be frequency-specific in the modelling of the causal relationship between signals Y and X, keeping the predicted value Y_t in the time domain. It allows us further to capture short- and long-range dependencies between signals with only few parameters to estimate and to

be relatively independent to the sampling frequency of the data.

6 Bayesian Granger-Causality Statistic

Based on the model proposed in this article, a necessary and sufficient condition for X not to be Granger-causal for Y_t at a given time t, is that each element of the subset $\widetilde{\varphi}_t$ of φ_t, that contains all the causal coefficients of interest, equals zero.

The most appropriate approach to evaluate the compatibility of this type of hypothesis with the data would be to compute for each time t a Bayes factor between the VAR model under the restriction $\widetilde{\varphi}_t = 0$ (M_1^t) and the VAR model without restrictions (M_2^t). This would requires the computation of the evidence of two models M_1^t and M_2^t, which becomes infeasible with our model (1).

It seems therefore reasonable in to rely on a simpler approach that is only based on the posterior density of the (unconditional) model. The use of highest posterior density (HPD) regions for Bayesian testing was introduced in Box and Tiao (1973) and used in the Bayesian literature, as for example in Kim and Press (1994) and West and Harrison (1997, p. 280). For a given time t, consider the sub-vector $\widetilde{\varphi}_t$ and let c be its dimension. A suitable partition of the posterior parameters $\mu_{\mathbf{t}}$ and Σ_t gives the marginal posterior density $q(\widetilde{\varphi}_t|Z_1^T) = \mathcal{N}_c(\widetilde{\mu}_t, \widetilde{\Sigma_t})$. The contribution of each element of $\widetilde{\varphi}_t$ to the prediction of Y_t may be assessed by considering the compatibility in this marginal posterior with the value $\widetilde{\varphi}_t = \widetilde{\varphi_{0_t}}$, where $\widetilde{\varphi_{0_t}}$ is a zero vector of dimension c. The question is therefore to know whether the parameter point $\widetilde{\varphi}_t = \widetilde{\varphi_{0_t}}$ is included in the highest posterior density region (HPD) of size $1-\alpha$. We thus need to calculate the probability of the event $p(\widetilde{\varphi}_t|Z_1^T) > p(\widetilde{\varphi_{0_t}}|Z_1^T)$. Now the particular point $\widetilde{\varphi_{0_t}}$ is included in the HPD region of content $1-\alpha$ if and only if $(\widetilde{\varphi_{0_t}} - \widetilde{\mu}_t)'\widetilde{\Sigma_t}^{-1}(\widetilde{\varphi_{0_t}} - \widetilde{\mu}_t) < k$, where k is the $1-\alpha$ quantile of the standard χ^2 distribution

with c degrees of freedom. As stated in Box and Tiao (1973, p. 125), it follows that the parameter point $\widetilde{\varphi_{0_t}}$ is covered by the HPD region of content $1-\alpha$ if and only if $\Pr\{p(\widetilde{\varphi}_t|Z_1^T) > p(\widetilde{\varphi_{0_t}}|Z_1^T)|Z_1^T\} \leq 1-\alpha$. Equivalently, we can search the region of minimum coverage that contains $\widetilde{\varphi_{0_t}}$.

7 Assessment of Accuracy

We now turn our attention to the accuracy of variational Bayesian inference for our model (1) under the mean-field assumption described in Section 3.1.2 and priors discussed in Section 4. We provide here a study of the quality of variational Bayesian estimate for model order selection and Granger causality detection. The first simulation study in Section 7.2 therefore provides an overview of the quality of the variational lower bound as a model order selection criterion and the second simulation study in Section 7.3 evaluates the ability of the BSS and MSBSS models to detect Granger causalities between two signals. We mention that a Monte-Carlo type of simulation was not feasible due to the very large amount of unknown parameters of our model and so direct comparison of the obtained variational posterior with true posterior density was unfortunately not feasible.

7.1 Practical implementation

To initialize the BSS and the MSBSS model algorithms, we run 10 iterations of the simple EM algorithm (Shumway and Stoffer, 1982). We thereby obtain reliable starting values for $\{\varphi_1, \Sigma_1, \Omega_1^b\}$ under initial conditions defined in Section 4. We then iteratively update the parameters of the variational approximate posteriors through the M-step equations described in Section 4.2.2 and the time-varying VAR coefficients through the modified Kalman smoothing E-step equations discussed in Section 4.2.1. These variational EM-

steps are iterated until the relative free-energy criterion described in Section 3.2 between two consecutive iterations changes less than a tolerance value that we choose here to be equal 10^{-2} for the model order selection study and to 10^{-4} for the Granger causality detection study following Menictas and Wand (2013).

7.2 Model order selection

The first set of simulations was run to assess the quality of the free energy as a model selection criterion. We simulated bivariate systems with numbers of time points equal to 500, 1000 and 2000 with 1 and 10 trials and model order 1, 2, 4 and 8. The data are simulated directly from the model (1). The matrices R and Q equal $I_d \times 0.1$ and $I_k \times 0.1$ respectively. The transition matrix of the state A was set to 0.85 for the diagonal element related to the real model order p, and decreasing values are set for the other diagonal elements $\{p - i\}_{i=1}^{p-1}$ (the values are $0.54, 0.34, 0.21, 0.13, 0.09, 0.05, 0.03$). The estimation procedure was performed for orders $p = 1, \ldots, 10$ and the order that exhibits the maximum free-energy quantity F_{m_p} was selected. The tolerance level between two iterations was chosen to be equal to 10^{-2} for computational simplicity. Estimation procedure for each model order was based on the same amount of data (i.e., conditioning on the first p observations), so that different models could be compared through their evidence (or free energy). For each setting, 100 replications were performed. Globally, the proportion of correctly selected models is very high (from 83% to 100%). This proportion grows up with the number of time points and the number of trials. These results show that the free-energy quantity acts as a very good model order selection criterion.

7.3 Granger-causality detection

We will now assess the ability of the proposed models BSS and MSBSS to detect Granger causality. For that purpose we simulated signals in order to observe how the method performs with data that are not generated according to the model, but with parameters that vary slowly in time, which is a reasonable simulation of neuroscientific data. The simulation consists in 50 replications of a bivariate BSS model but with hidden variables φ_1^T that evolve slowly and deterministically through time. Based on these φ_1^T, data are generated from the observed equation in model (1) with deterministic parameters φ_1^T. The simulations were carried out for model orders $\{1; 2; 4; 8\}$, series lengths $\{500; 1000; 2000\}$, number of trials $\{1; 10\}$ and also for different values of the causal parameter $\{1; 0.8; 0.6; 0.4; 0.2\}$, to see the limits of the methods. The simulations for the causal parameter values $\{0.8; 0.6; 0.4; 0.2\}$ were carried out for all model orders and number of trials, but only for a series length of 500 time points, which is the case where the method will break first.

The R matrix was fixed to the $I_d \times 0.1$ values and the slowly-varying parameters are all set to zero except for the entries related to the causal parameters $\{\varphi_{2\to1}\}_{t=1}^T$ and $\{\varphi_{1\to1}\}_{t=1}^T$ for order p. The values of the simulated parameters can be seen on the two panels (which are identical) on the top of Figures in Appendix M. Signals according to slowly-varying parameters were simulated with normal and non-normal errors.

Data generated with slowly-varying parameters and normal errors. Concerning the the MSBSS model, for each of the 50 simulation, the model order p_j for the different scales as well as the number of scales J to be taken in the model are selected by comparing the free-energy quantities as described in Section 5. The maximum number of scale was set at 4 and the maximal model order per scale at 5. In multiple trial scenarios, the number of scales as well as the model order selection procedure was carried on one single trial only for

computational simplicity. Estimation procedure is always conditioned on the same number of time points, allowing us to compare the models. The MSBSS model is then estimated and the detection capability of the Granger causality is calculated. The Granger-causality detection capability of a method is quantified by the true negative rate (TNR) and the true positive rate (TPR). For each time, TNR is the percentage of the 50 simulations for which the 95% HPD region defined in Section 6 contains the causal statistic when it must actually contain it, and TPR is the percentage of the 50 simulations for which the 95% HPD region does not contain the causal statistic when it should not. The significance assessment is based on the χ^2 threshold as explained in Section 7.3. Secondly, the BSS model is estimated with the true model order (oracle order) and the relative TNR and TPR are computed for that model. Finally, for the BSS model, the model order p is selected with the free-energy criterion as described in Section 3.2 and the model is estimated based on this selected order p. The relative TNR and TPR are calculated.

The causality detection ability for the BSS model with estimated model order is not shown in the results, because the model order was selected correctly for each simulation. The results are thus the same as those for the BSS model with oracle order.

Figure 2 shows Granger-causality detection accuracy for a model order 1, series length of 500 and causal parameter equal to 1. By construction, the causality arises between times $190 - 500$, as shown in the two graphs on the top of Figure 2. Here, the MSBSS and BSS models yield very similar results, correctly recovering the underlying directional influences with 1 and even better with 10 trials. One can observe that the causality detection evolves as the causal parameter changes from zero to one. The TPR is indeed gradually increasing as the parameter changes from zero to one, slightly faster for the BSS than for the MSBSS model for $N = 1$, and identically for $N = 10$.

Figure 3 shows Granger-causality detection accuracy for a model order 1, series length of

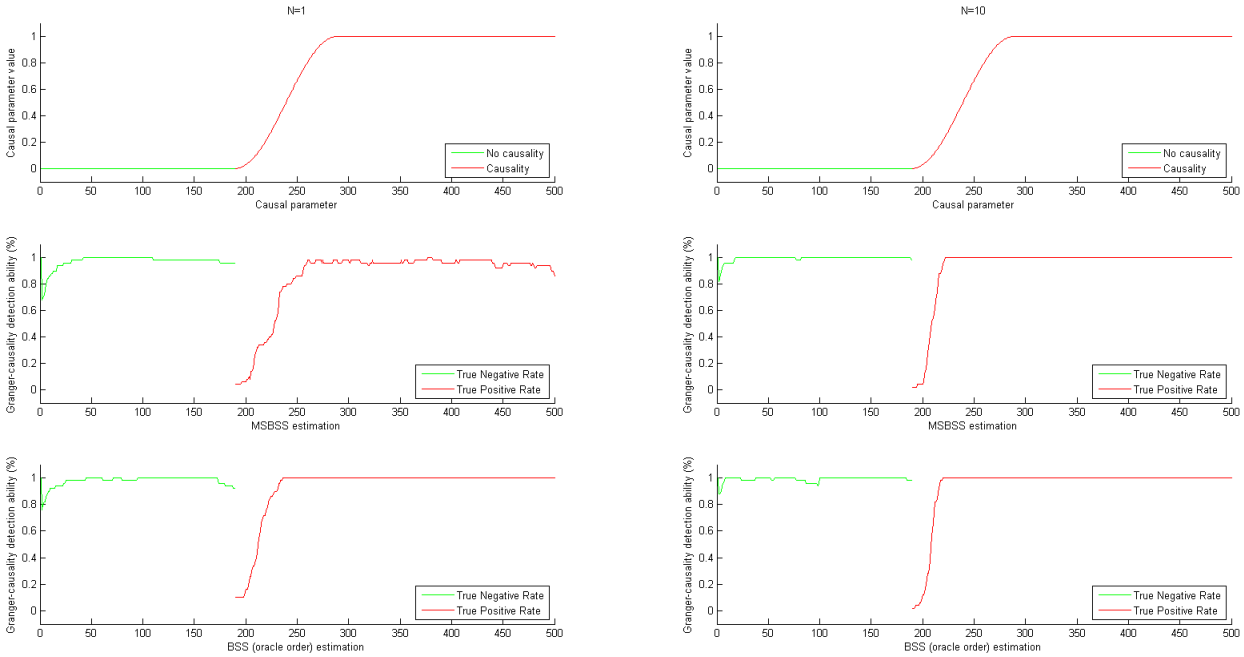

Figure 2: Granger-causality detection ability for order 1, series length 500 and causal parameter 1. Top graphs show the value of the true parameters. Middle graphs display the true positive and true negative rates for the MSBSS model and bottom graphs for the BSS model with the true model order (oracular). This figure appears in color in the electronic version of this article.

500 and causal parameter 0.2. By construction, the causality arises between times $190-500$ as shown in the two graphs on the top of Figure 3. For such a low causal parameter, the TPR related to the MSBSS and the BSS models for 1 trial is null, but surprisingly high for 10 trials, around 70% for the MSBSS model and around 90% for the BSS model.

Figures for Granger-causality detection accuracy related to all other model orders, series

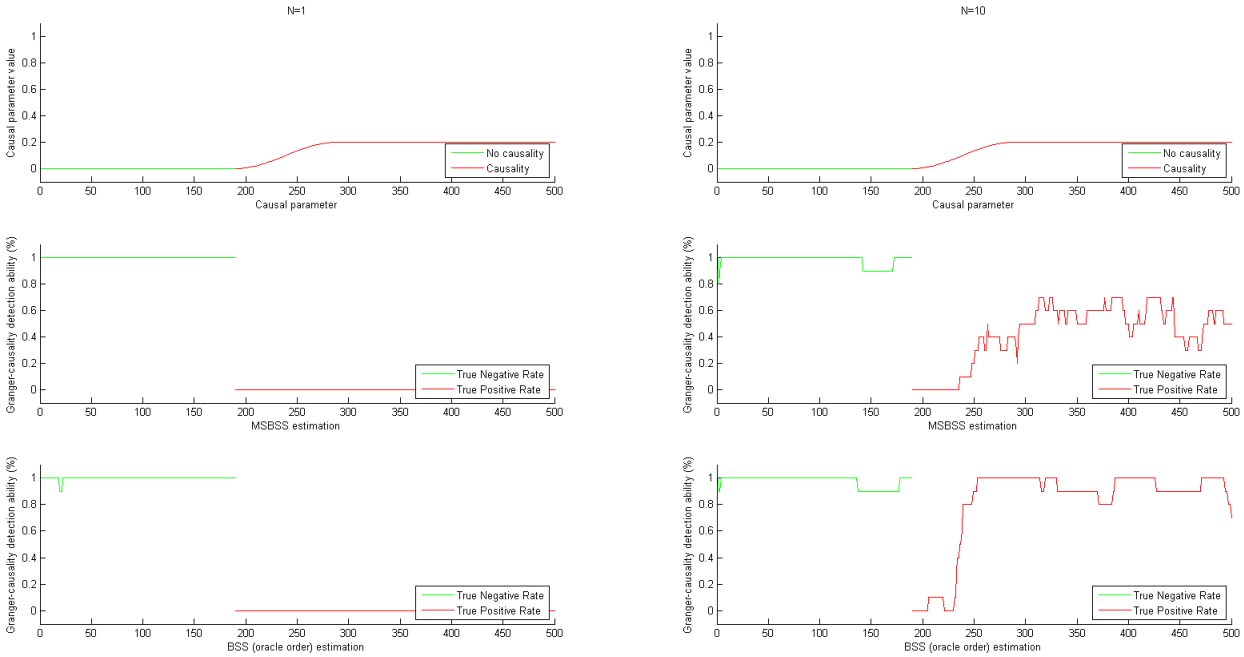

Figure 3: Granger-causality detection ability for order 1, series length 500 and causal parameter 0.2. See Figure 2 for more details. This figure appears in color in the electronic version of this article.

lengths and causal parameter values for normal errors are given in Appendix M. Globally, for a causal parameter that equals 1, the MSBSS and the BSS models show very good Granger-causality detection accuracy in terms of TPR and TNR for 1 and 10 trials. The only cases which display poorer results are the causality detection for the MSBSS model with 1 trial for model order 2 and series length 500, 1000 and 2000. When the causal parameter value decreases, the causality detection ability decreases as well. Poorer causality detection results are globally observed for a causal parameter from 0.6 to 0.2 for 1 trial and

for a causal parameter value of 0.2 for 10 trials, especially for the MSBSS model.

Data generated with slowly-varying parameters and non-normal errors. Data with slowly-varying parameters and non-normal errors are generated as explained in Section 7.3, with the errors following a multivariate t-distribution with parameters $\nu = 5$ and $R = I_d \times 0.1$. The observation equation errors are therefore multivariate t-distributed. The simulations were carried out for different model orders $\{1; 2; 4; 8\}$, a series length of 500 time points, 1 and 10 trials and for different values of the causal parameter $\{1; 0.8; 0.6; 0.4; 0.2\}$.

All graphical results related to the data generated with slowly-varying parameters and non-normal errors can be found in Appendix M. Globally, the detection accuracy remains satisfactory with data generated with non-normal errors. The Granger-causality detection ability of the proposed methods is therefore robust to this model assumption departure.

8 Application

We will now apply the proposed MSBSS model to real iEEG (intracranial electroencephalogram or local field potential) data recorded during a psychological experimental situation. Brain recordings, here in one patient suffering of a pharmaco-resistant epilepsy and candidate for a selective brain resection, are localized within the amygdala (AMY) and medial orbito-frontal cortex (mOFC) regions in order to study the dynamics of neuronal processes within and between these regions in response to emotional prosody exposure. It is known in the literature that the emotional content of the stimulus induces the presence of causal links AMY $\rightarrow$ mOFC and mOFC $\rightarrow$ AMY.

8.1 Results

Here we present the results for one patient and the experimental condition anger. The patient was exposed to short pseudowords pronounced with an angry prosody (ADD REF Grandjean et al., 2005). The data were acquired with a high resolution of 512Hz during 2.25sec. There are 28 trials available. As high frequency behaviour of data was not of interest, we downsampled the data by a factor of four with an exponential smoother (it is of course vital to process with a smoother that is not based on future values). A MSBSS model was estimated following the procedure described in Section 5. As researchers were already aware about frequency bands where they expect a causal relationship, the number of scales $J_{\max}$ was fixed to 4 (on the downsampled data), corresponding to four frequency bands respectively around 64, 32, 16 and 8Hz, plus a smooth that represents the frequency content below 8Hz. The selection of the maximal model order per scale p_j was done independently for each scale, with a maximal value of $p_{\max} = 5$ for the scales related to the frequency bands 64Hz and 32Hz, $p_{\max} = 3$ for the scale related to the frequency band 16Hz and $p_{\max} = 1$ for the scales related to the frequency band 8Hz and the smooth. These particular choices for $p_{\max}$ are motivated by the occurrence of the onset of the stimulus at 250ms. Indeed, due to the use of past values in the model, the estimation procedure starts at time $\max p_{\max,j} 2^j + 1$ and we do not want to miss the onset of the stimulus in the model estimation. The model order per scale p_j is selected on the whole set of trials.

Based on the MSBSS model, a significant Granger causality from signal 1 to signal 2 at frequency f and at time t means that the energy of signal 1 at frequency f significantly improves the prediction of the value of signal 2 at time t.

The testing procedure was performed as follows: we first tested the overall set of causal VAR parameters of interest and then the scale-specific causal VAR parameters ($\widetilde{\varphi}_t$) by considering the compatibility in their marginal posterior with the value $\widetilde{\varphi}_t = \widetilde{\varphi_{0_t}}$, where

$\widetilde{\varphi_{0_t}}$ is a zero vector of dimension c. For this purpose we computed the highest posterior density (HPD) region, as in Section 6. Instead of fixing a content of coverage (typically 95%) and determining if $\widetilde{\varphi_{0_t}}$ is included in the HPD region of that content, we computed the coverage of the smallest HPD region that contains $\widetilde{\varphi_{0_t}}$. This gives us more information and we will call the complementary probability of this coverage the significance level.

The results are reported for the two directional causalities AMY $\rightarrow$ mOFC and mOFC $\rightarrow$ AMY. The first line of each graph in Figure 4 represents the results related to the overall statistic and each following line represents the results related to each scale corresponding to the frequency bands respectively around 64, 32, 16 and 8Hz.

The number of tests provided by the proposed method is basically proportional to the number of time points. To circumvent the multiple testing problem, a solution that seems suitable is the cluster mass test which consists in defining clusters of neighbouring time regions (Maris and Oostenveld, 2007). A cluster is defined as the sum of the statistics in a segment where each individual statistic (i.e. the time-dependent Granger-causality statistic) is above a given threshold. Here we chose a threshold corresponding to a level of .2 (this choice does not affect the familywise error rate and minimally affect the sensitivity of the test). The different trials were then permuted between brain regions AMY and mOFC in order to give new random partitions and the full procedure were applied for each new partition. The largest cluster statistic was then retained for each partition and this procedure was repeated 2000 times in order to get a reliable distribution of the largest cluster statistic under the null hypothesis of Granger non-causality.

8.1.1 Testing the scale specific causality

Figure 4 shows the results of the estimated Granger causality for the two causal links of interest AMY $\rightarrow$ mOFC and mOFC $\rightarrow$ AMY. Vertical lines represent the onset and the

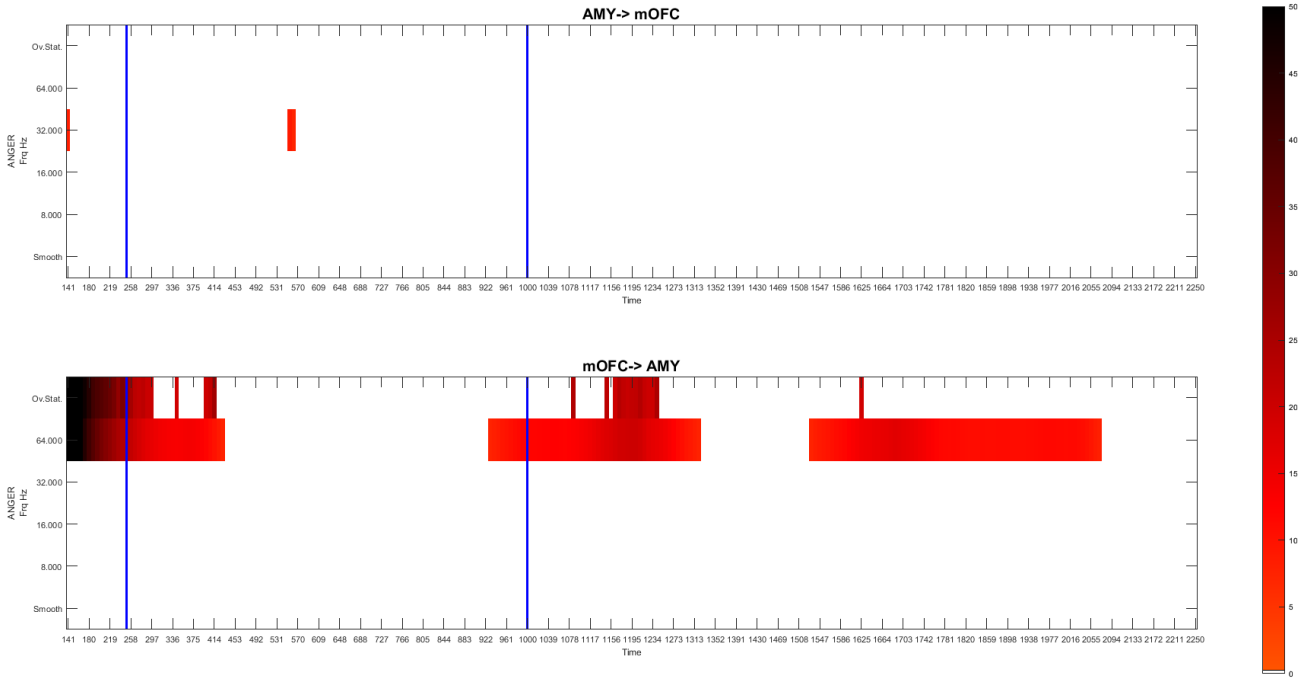

Figure 4: Results of the estimated Granger causality for the two causal links of interest AMY $\rightarrow$ mOFC and mOFC $\rightarrow$ AMY for the overall causality statistic and the scale specific causality statistics. Intensity represents the value of the individual statistic. Only the clusters assessed significant by the cluster mass test are displayed. Vertical lines represent the onset and the offset of the stimulus which occur respectively at 250ms and 1000ms. This figure appears in color in the electronic version of this article.

offset of the stimulus which occur respectively at 250ms and 1000ms. The estimated model order for each scale plus the smooth is [5 5 3 1 1]. Only the clusters assessed significant by the cluster mass test are displayed and the colors represent the value of the individual statistic. As mentioned in Section 5, the smooth represents the frequency content of the

series from the largest scale to the lowest frequency in the signal.

These results give a partial answer to the question of how AMY and mOFC regions are functionally and causally related during the exposure of auditory emotional stimuli. Actually, an initial Granger causality event from mOFC $\rightarrow$ AMY is observed in gamma range (64Hz) just before the stimuli onset and during stimulus exposure followed by an AMY $\rightarrow$ mOFC in beta range (32Hz) during this same period. At the offset of the stimuli the Granger causality is again from mOFC $\rightarrow$ AMY in gamma range (64Hz), this directional Granger causality is sustained during the period after the stimuli presentation. These results are compatible with a known complex cross-talk between these two brain regions during emotional stimulus exposure and the related meaning of such event for the organism. Of course such results should be extended to several patients with similar brain recordings in targeted brain regions and compared to the processing of neutral stimuli.

9 Conclusion

We derived a time-varying Granger-causality statistic through a Bayesian nonstationary multivariate time series model with dynamic coefficients. While similar models have been proposed (Cassidy, 2002), one of the main contributions of this article is to provide an assessment of accuracy of the method, and especially an extension to the *à trous* Haar wavelets transform (Renaud et al., 2003). This very flexible *à trous* Haar procedure enables us to capture short- and long-range dependencies between signals with only few parameters to be estimated. The central finding of this article is that variational Bayesian estimate of time-varying VAR models enables us to achieve good (if not excellent) accuracy in terms of Granger-causality detection for normal, but also for non-normal errors. The simulation study also illustrates that the free-energy criterion enables us to select the model order

with an high degree of accuracy. A further point to highlight is the suitability of the model to deal with multiple trial datasets in a fully correct way. Each epoch is indeed treated as a separate observation sequence and all of them are taken properly into account in the model and in the estimation procedure.

From the experimental neuroscientific point of view, the method has a number of desirable properties. It performs very well in terms of Granger-causality detection. With the proposed *à trous* Haar extension, the MSBSS model enables us to be specific for the frequency in the assessment of causality, which is a main point of interest in the neuroscience community. We show that the MSBSS model is able to capture efficiently the causality pattern between two signals, due to its ability to take into account the short- and long-range dependence components. Thus, the variational Bayesian framework combined with the *à trous* Haar methodology for modelling time-varying multivariate time series provides a potentially very powerful tool for dynamical spectral causality analysis in a neuroscientific context.

SUPPLEMENTARY MATERIAL

1 Introduction **3**

2 Granger causality **4**

3 The Bayesian State Space Model **5**

3.1 Elements of variational Bayes 6

3.1.1 Learning rules 7

3.1.2 Mean-field approximation 7

3.1.3 Variational EM algorithm 7

3.1.4 Conjugate exponential 8

3.2 The variational evidence lower bound 9

4 Model specification **9**

4.1 Prior distributions 9

4.1.1 Prior for φ_1^T 9

4.1.2 Prior for $\Omega_1 = \{A, \alpha, \delta\}$ 11

4.1.3 Prior for $\Omega_2 = \{Q, a_q\}$ 12

4.1.4 Prior for $\Omega_3 = \{R, a_r\}$ 13

4.2 Update equations 13

4.2.1 Variational E-step 14

4.2.2 Variational M-step 14

4.3 Multiple trials 16

5 The Multiscale Bayesian State Space Model **16**

5.1 The *à trous* Haar wavelets transform 17
5.2 The multiscale Bayesian state space model 17

6 Bayesian Granger-Causality Statistic 19

7 Assessment of Accuracy 20
7.1 Practical implementation 20
7.2 Model order selection 21
7.3 Granger-causality detection 22

8 Application 26
8.1 Results 27
8.1.1 Testing the scale specific causality 28

9 Conclusion 30

A Appendix A: Computation of the Free Energy 37

B Appendix B: Mean and Fluctuation Theorem 38

C Appendix C: Unified Inference Theorem 39

D Appendix D: E-step: Computation of the Distribution of φ_1^T 42

E Appendix E: M-step: Computation of the Distribution of δ 44

F Appendix F: M-step: Computation of the Distribution of α 45

G Appendix G: M-step: Computation of the Distribution of A 47

H Appendix H: M-step: Computation of the Distribution of a_q **49**

I Appendix I: M-step: Computation of the Distribution of Q **50**

J Appendix J: M-step: Computation of the Distribution of a_r **52**

K Appendix K: M-step: Computation of the Distribution of R **53**

L Appendix L: Multiple trials **54**

M Appendix M: Granger-Causality Detection Results **59**

M.1 Data generated with slowly-varying parameters and normal errors 59

M.2 Data generated with slowly-varying parameters and non-normal errors 84

References

Barber, D. and S. Chiappa (2007). Unified Inference for Variational Bayesian Linear Gaussian State-Space Models. In B.Schölkopf, J.C.Platt, and T.Hofmann (Eds.), *Advances in Neural Information Processing Systems 19 (NIPS 2006)*, pp. 81–88. MIT Press.

Beal, M. (2003). *Variational Algorithms for Approximate Bayesian Inference*. Ph. D. thesis, Gatsby Computational Neuroscience Unit, University College London, London.

Box, G. E. P. and G. C. Tiao (1973). *Bayesian Inference in Statistical Analysis*. Addison-Wesley, Reading, Massachusetts.

Cassidy, M.J. Penny, W. (2002). Bayesian nonstationary autoregressive models for biomedical signal analysis. *IEEE Transactions on Biomedical Engineering 49*(10), 1142–1152.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis 1*(3), 515–534.

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica 37*(3), 424–438.

Huang, A., M. P. Wand, et al. (2013). Simple marginally noninformative prior distributions for covariance matrices. *Bayesian Analysis 8*(2), 439–452.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering 82*(1), 35–45.

Kalman, R. E. and R. S. Bucy (1961). New Results in Linear Filtering and Prediction Theory. *Journal of Basic Engineering 83*(1), 95–108.

Kim, H. J. and S. J. Press (1994). *Bayesian hypothesis testing of equality of normal covariance matrices*, Volume 24 of *Lecture Notes–Monograph Series*, pp. 323–330. Hayward, CA: Institute of Mathematical Statistics.

Lütkepohl, H. (2005). *New Introduction to Multiple Time Series Analysis*. Cambridge University Press.

Maris, E. and R. Oostenveld (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods 164*(1), 177–190.

Menictas, M. and M. P. Wand (2013). Variational inference for marginal longitudinal semiparametric regression. *Stat 2*(1), 61–71.

Ostwald, D., E. Kirilina, L. Starke, and F. Blankenburg (2014). A tutorial on variational Bayes for latent linear stochastic time-series models. *Journal of Mathematical Psychology 60*, 1–19.

Renaud, O., J.-L. Starck, and F. Murtagh (2003). Prediction Based on a Multiscale Decomposition. *International Journal of Wavelets, Multiresolution and Information Processing 1*(2), 217–232.

Shumway, R. H. and D. S. Stoffer (1982). An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis 3*(4), 253–264.

Wand, M. P., J. T. Ormerod, S. A. Padoan, and R. Fuhrwirth (2011). Mean Field Variational Bayes for Elaborate Distributions. *Bayesian Analysis 6*(4), 847–900.

West, M. and P. J. Harrison (1997). *Bayesian Forecasting & Dynamic Models*. New York: Springer.

A Appendix A: Computation of the Free Energy

It is straightforward to re-express the free-energy quantity as

$$\begin{aligned}\mathbf{F} = &- \left\langle \log \frac{q(\Omega_1^b|Z_1^T)}{p(\Omega_1^b)} \right\rangle_{q(\Omega_1^b|Z_1^T)} - \left\langle \log q(\varphi_1^T|Z_1^T) \right\rangle_{q(\varphi_1^T|Z_1^T)} \\ &+ \left\langle \log p(\varphi_1^T, Z_1^T|\Omega_1^b, m_p) \right\rangle_{q(\varphi_1^T,\Omega_1^b|Z_1^T)},\end{aligned} \tag{13}$$

where we omit the conditional dependence to the model m_p for notational simplicity. The first term on the r.h.s. of equation (13) represents the Kullback–Leibler divergence between the prior and the variational posterior for the entire set of parameters Ω_1^b, the second r.h.s. term represents the entropy of the variational posterior distribution of the hidden variables φ_1^T and the third r.h.s. term is the average log-likelihood of the data and hidden state parameters taken over the entire set of parameters Ω_1^b and hidden state φ_1^T.

Recalling the mean-field factorization assumed in our model as well as the priors distributions set in Section 4, the free-energy quantity takes the explicit form:

$$\begin{aligned}\mathbf{F} = &- \left\langle \log \frac{q(A)}{p(A|\alpha)} \right\rangle_{q(A)} - \left\langle \log \frac{q(\alpha|\delta)}{p(\alpha|\delta} \right\rangle_{q(\alpha|\delta)} - \left\langle \log \frac{q(\delta)}{p(\delta)} \right\rangle_{q(\delta)} \\ &- \left\langle \log \frac{q(Q|a_q)}{p(Q|a_q)} \right\rangle_{q(Q|a_q)} - \left\langle \log \frac{q(a_q)}{p(a_q)} \right\rangle_{q(a_q)} \\ &\underbrace{- \left\langle \log \frac{q(R|\{a_{ri}\}_{i=1}^d)}{p(R|\{a_{ri}\}_{i=1}^d)} \right\rangle_{q(R|\{a_{ri}\}_{i=1}^d)} - \left\langle \log \frac{q(\{a_{ri}\}_{i=1}^d)}{p(\{a_{ri}\}_{i=1}^d)} \right\rangle_{q(\{a_{ri}\}_{i=1}^d)}}_{\text{term 1}} \\ &\underbrace{- \left\langle \log q(\varphi_1^T) \right\rangle_{q(\varphi_1^T)}}_{\text{term 2}} \underbrace{+ \left\langle \log p(\varphi_1^T, Z_1^T|\Omega_1^b) \right\rangle_{q(\varphi_1^T)q(\Omega_1^b)}}_{\text{term 3}},\end{aligned} \tag{14}$$

where the conditional dependance of the variational posterior distributions $q(.)$ to the data Z_t is now omitted for notational simplicity.

The KL divergences of term 1 have closed form for conjugate exponential distributions and are therefore straightforward to obtain despite their intensive computation. Let us

focus on the entropy term that appears in term 2. Following Beal (2003), this can be rewriten as

$$\text{term 2} = -\left\langle \log q(\varphi_1^T) \right\rangle_{q(\varphi_1^T)} = -\left\langle -\log \Upsilon + \underbrace{\left\langle \log p(\varphi_1^T, Z_1^T | \Omega_1^b) \right\rangle_{q(\Omega_1^b)}}_{\text{term 3}} \right\rangle_{q(\varphi_1^T)} . \tag{15}$$

Term 3 therefore disappears in equation (14) and the quantity Υ becomes

$$\begin{aligned} \Upsilon &= \left\langle \exp \left\langle \log p(\varphi_1^T, Z_1^T | \Omega_1^b) \right\rangle_{q(\Omega_1^b)} \right\rangle_{q(\varphi_1^T)} \\ &= \underbrace{\left\langle \exp \left\langle \log p(\varphi_1^T | Z_1^T, \Omega_1^b) \right\rangle_{q(\Omega_1^b)} \right\rangle_{q(\varphi_1^T)}}_{=1} + \left\langle \exp \left\langle \log p(Z_1^T | \Omega_1^b) \right\rangle_{q(\Omega_1^b)} \right\rangle_{q(\varphi_1^T)} . \end{aligned} \tag{16}$$

Assuming now that the parameters Ω_1^b have a point mass density rather than their variational posterior distribution $q(.)$, the Υ quantity becomes

$$\Upsilon = \left\langle \exp \left\langle \log p(Z_1^T | \Omega_1^b) \right\rangle_{q(\Omega_1^b)} \right\rangle_{q(\varphi_1^T)} = p(Z_1^T | \bar{R}) = p(Z_1) \prod_{t=2}^{T} p(Z_t | Z_1^{t-1}). \tag{17}$$

The quantity Υ can be obtained just after the forward recursion step and therefore amounts to

$$\Upsilon = p(Z_1) \prod_{t=2}^{T} p(Z_t | Z_1^{t-1}) = \prod_{t=1}^{T} \mathcal{N}_d(C_t \mu_t^{t-1}, C_t \Sigma_t^{t-1} C_t' + R), \tag{18}$$

where the quantities μ_t^{t-1} and Σ_t^{t-1} are defined in Appendix D (**?**).

B Appendix B: Mean and Fluctuation Theorem

The *Mean and Fluctuation Theorem* is a decomposition proved by Barber and Chiappa (2007). Using our notation, based on the model defined by equations (1) and the Conditional distributions in Section 3 and the set of unknown parameters Ω_1^b, the following

decomposition of the density $\left\langle \log p(\varphi_1^T, Z_1^T|\Omega_1^b)\right\rangle_{q(\Omega_1^b|Z_1^T)}$ holds

$$\begin{aligned}
\left\langle \log p(\varphi_1^T, Z_1^T|\Omega_1^b)\right\rangle_{q(\Omega_1^b|Z_1^T)} =& \log p(\varphi_1^T, Z_1^T|\left\langle \Omega_1^b\right\rangle) + F_{A,Q} + F_{C_t,R}, \quad \text{where} \\
\log p(\varphi_1^T, Z_1^T|\left\langle \Omega_1^b\right\rangle) \propto& \frac{1}{2}(\varphi_1 - \mu_1)\Sigma_1^{-1}(\varphi_1 - \mu_1)' \\
&- \frac{1}{2}\sum_{t=2}^{T}(\varphi_t - \langle A\rangle \varphi_{t-1})Q^{-1}(\varphi_t - \langle A\rangle \varphi_{t-1})' \\
&- \frac{1}{2}\sum_{t=1}^{T}(Z_t - C_t\varphi_t)R^{-1}(Z_t - C_t\varphi_t)', \qquad (19) \\
F_{A,Q} =& -\frac{1}{2}\sum_{t=1}^{T-1}\varphi_t'\big(\left\langle AQ^{-1}A\right\rangle - \langle A\rangle \langle Q\rangle^{-1}\langle A\rangle\big)\varphi_t, \\
F_{C_t,R} =& -\frac{1}{2}\sum_{t=1}^{T}\varphi_t'\big(\left\langle C_tR^{-1}C_t\right\rangle - \langle C_t\rangle \langle R\rangle^{-1}\langle C_t\rangle\big)\varphi_t.
\end{aligned}$$

Note that the formulation in Barber and Chiappa (2007) considers the matrix C_t as time-invariant.

C Appendix C: Unified Inference Theorem

For the BSS model, recall equation (5), which is a key quantity we want to evaluate

$$q^*(\varphi_1^T)^{(t+1)} \propto \exp\left\langle \log p(\varphi_1^T, \Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b)}^{(t)} \propto \exp\left\langle \log p(\varphi_1^T|\Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b)}^{(t)}.$$

In the situation where Ω_1^b are random parameters rather than fixed values, Barber and Chiappa (2007) propose an elegant solution based on a suitably augmented system of equations that allows to infer $q(\varphi_1^T)$ through classical state space model inference algorithms. The theorem states that the above density can be written as

$$\exp\left\langle \log p(\varphi_1^T|\Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b)} = p(\varphi_1^T|\widetilde{\Omega_1^b}, \widetilde{Z}_1^T), \qquad (20)$$

where the augmented elements are

$$\widetilde{\Omega}_1^b = \{\widetilde{A}; \widetilde{Q}; \widetilde{R}\}, \qquad \widetilde{A} = \langle A \rangle, \qquad \widetilde{Q} = \langle Q \rangle, \qquad \widetilde{R} = \begin{pmatrix} \langle R \rangle & 0 & 0 \\ 0 & I_k & 0 \\ 0 & 0 & I_d \end{pmatrix},$$

$$\widetilde{Z}_t = \begin{pmatrix} Z_t \\ 0_k \\ 0_d \end{pmatrix}, \qquad \widetilde{C}_t = \begin{pmatrix} C_t \\ U_A \\ U_{C_t} \end{pmatrix},$$

and where U_A and U_{C_t} are defined as the Cholesky decompositions of

$$U_A' U_A = \left\langle A Q^{-1} A \right\rangle - \langle A \rangle \langle Q \rangle^{-1} \langle A \rangle,$$
$$U_{C_t}' U_{C_t} = \left\langle C_t R^{-1} C_t \right\rangle - \langle C_t \rangle \langle R \rangle^{-1} \langle C_t \rangle.$$

In our specific case, we have that $U_{C_t}' U_{C_t} = 0$, due to the non-randomness of the matrix C_t, and so the unique quantity to define is $U_A' U_A$. In the case where A and Q are diagonal

and defined respectively by equations (7) and (9), we have that

$$
\begin{aligned}
U_A' U_A &= \left\langle AQ^{-1}A \right\rangle - \langle A \rangle \langle Q \rangle^{-1} \langle A \rangle \\
&= \begin{pmatrix} \left\langle A_{1,1} q_1^{-1} A_{1,1} \right\rangle & 0 & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & \left\langle A_{k,k} q_k^{-1} A_{k,k} \right\rangle \end{pmatrix} \\
&\quad - \begin{pmatrix} \langle A_{1,1} \rangle \left\langle q_1^{-1} \right\rangle \langle A_{1,1} \rangle & 0 & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & \langle A_{k,k} \rangle \left\langle q_k^{-1} \right\rangle \langle A_{k,k} \rangle \end{pmatrix}, \\
&= \begin{pmatrix} \left\langle A_{1,1}^2 \right\rangle \left\langle q_1^{-1} \right\rangle & 0 & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & \left\langle A_{k,k}^2 \right\rangle \left\langle q_k^{-1} \right\rangle \end{pmatrix} - \begin{pmatrix} \langle A_{1,1} \rangle^2 \left\langle q_1^{-1} \right\rangle & 0 & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & \langle A_{k,k} \rangle^2 \left\langle q_k^{-1} \right\rangle \end{pmatrix}, \\
&= \begin{pmatrix} \sigma_{A_{1,1}}^2 \left\langle q_1^{-1} \right\rangle & 0 & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & \sigma_{A_{k,k}}^2 \left\langle q_k^{-1} \right\rangle \end{pmatrix},
\end{aligned}
\tag{21}
$$

where $\sigma_{A_{i,i}}^2$ is the variational posterior variance of the i-th entry of the A matrix defined in Section 4.2.2 and $\left\langle q_i^{-1} \right\rangle$ is the variational posterior mean of the i-th entry of the inverse variance-covariance matrix Q^{-1} defined in Section 4.2.2 that can be straightforwardly obtained due to the properties of the inverse-gamma and gamma distributions.

For our model, the augmented system of equations then gives

$$
\widetilde{Z}_t = \begin{pmatrix} Z_t \\ 0_k \end{pmatrix}, \qquad \widetilde{A} = \langle A \rangle, \qquad \widetilde{Q} = \langle Q \rangle, \qquad \widetilde{C}_t = \begin{pmatrix} C_t \\ U_A \end{pmatrix}, \qquad \widetilde{R} = \begin{pmatrix} \langle R \rangle & 0 \\ 0 & I_k \end{pmatrix}, \tag{22}
$$

where U_A is defined in equation (21). Complete proofs can be found in Barber and Chiappa (2007) and Ostwald et al. (2014).

D Appendix D: E-step: Computation of the Distribution of φ_1^T

As discussed in Section 3.1, variational Bayesian algorithms lead to EM-like iterative equations for the optimal variational densities $q^*(.|Z_1^T)$. The variational update equation form for the hidden variables φ_1^T was state in equation (5). The derivation of equation (5) can be found in Appendix D. Due to the conjugacy condition discussed in Section 3.1.4 and equation the Conditional distributions stated in Section 3, the variational posterior distribution $q^*(\varphi_1^T|Z_1^T)$ is multivariate Gaussian of dimension $k \times T$. As explained in Beal (2003) and Cassidy (2002), if the parameters Ω_1^b were, as they called, point estimated, equation (5) would be straightforwardly resolved with classical state space model tools.

However in our variational Bayesian scenario, parameters Ω_1^b are random variables governed by a specific variational distribution $q(.|Z_1^T)$, and so $\exp\left\langle \log p(\varphi_1^T|\Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b|Z_1^T)}$ has to be computed for each variable with respect to its variational posterior distribution $q(\Omega_1^b|Z_1^T)$, sequentially for all parameters in the set Ω_1^b. As described in Beal (2003) and Barber and Chiappa (2007) and fully explained in Ostwald et al. (2014), the simplification in equation (5) no longer holds in this random parameter scenario because

$$\exp\left\langle \log p(\varphi_1^T|\Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b|Z_1^T)} \neq p(\varphi_1^T|\bar{\Omega}_1^b, Z_1^T). \tag{23}$$

The difference between the two terms in equation (23) is evaluated in Barber and Chiappa (2007) and its decomposition is known as the mean and fluctuation theorem, reproduced in Appendix B. To nevertheless use standard state space model algorithms to solve $\exp\left\langle \log p(\varphi_1^T|\Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b|Z_1^T)}$ in the variational Bayesian framework, and therefore to capitalize on the vast literature and results that exist on the topic, Barber and Chiappa (2007) prove the so-called unified inference theorem recalled in Appendix C. The

idea is to reformulate $\exp\left\langle \log p(\varphi_1^T, \Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b|Z_1^T)}$ as a standard density $\widetilde{p}(\varphi_1^T, \widetilde{\Omega}_1^b, \widetilde{Z}_1^T)$ with known parameters $\widetilde{\Omega}_1^b$, by suitably augmenting the first equation in the system (1). The classical KRTS smoother algorithms may then be suitably applied to the new density $\widetilde{p}(\varphi_1^T, \widetilde{\Omega}_1^b, \widetilde{Z}_1^T)$, and allows us to derive the target quantity $\widetilde{p}(\varphi_1^T|\widetilde{\Omega}_1^b, \widetilde{Z}_1^T)$. This finally corresponds to $\exp\left\langle \log p(\varphi_1^T|\Omega_1^b, Z_1^T)\right\rangle_{q(\Omega_1^b|Z_1^T)}$, and by extension to $q^*(\varphi_1^T|Z_1^T)^{(t+1)}$ by equation (5).

The optimal variational posterior distribution for the hidden state sequence φ_1^T is therefore multivariate Gaussian at each time t:

$$q^*(\varphi_t|Z_1^T) = \mathcal{N}_k(\varphi_t|\mu_t, \Sigma_t), \tag{24}$$

where the sufficient statistics $\{\mu_t; \Sigma_t\}_{t=1}^T$, as well as the cross-moments $\{\mu_t\mu_{t-1}; \Sigma_{t,t-1}\}_{t=2}^T$, are obtained through the KRTS smoother recursive equations applied to the suitable augmented system of equations discussed above.

Forward recursions. This step implements the recursive equation for $q(\varphi_t|Z_1^t)^{(l+1)}$. Let $\mu_t^t = \mathrm{E}_q\{\varphi_t|Z_1^t\}$, $\Sigma_t^t = \mathrm{VAR}_q\{\varphi_t|Z_1^t\}$, $\mu_t^{t-1} = \mathrm{E}_q\{\varphi_t|Z_1^{t-1}\}$, $\Sigma_t^{t-1} = \mathrm{VAR}_q\{\varphi_t|Z_1^{t-1}\}$ and $\langle . \rangle$ denotes the expectation with respect to the suitable variational distribution $q(.|Z_1^T)^{(l)}$. These quantities are obtained recursively as

$$\begin{aligned}
\mu_t^{t-1} &= \langle A \rangle \, \mu_{t-1}^{t-1}, \\
\Sigma_t^{t-1} &= \langle A \rangle \, \Sigma_{t-1}^{t-1} \, \langle A \rangle^T + \langle Q \rangle \,, \\
\mu_t^t &= \mu_t^{t-1} + K_t(Z_t - C_t\mu_t^{t-1}), \\
\Sigma_t^t &= \Sigma_t^{t-1} - K_t C_t \Sigma_t^{t-1},
\end{aligned} \tag{25}$$

where the Kalman gain K_t is given by

$$K_t = \Sigma_t^{t-1} C_t' (C_t \Sigma_t^{t-1} C_t' + \langle R \rangle)^{-1}. \tag{26}$$

Backward recursions. The backward recursions implement the recursive equations for $q(\varphi_t|Z_1^T)^{(l+1)}$. Let now $\mu_t^T = \mathrm{E}_q\{\varphi_t|Z_1^T\}$ and $\Sigma_t^T = \mathrm{VAR}_q\{\varphi_t|Z_1^T\}$. These quantities are obtained recursively as

$$\begin{aligned} J_t &= \Sigma_t^t \left\langle A \right\rangle' (\Sigma_{t+1}^t)^{-1}, \\ \mu_t^T &= \mu_t^t + J_t(\mu_{t+1}^T - \left\langle A \right\rangle \mu_t^t), \\ \Sigma_t^T &= \Sigma_t^t + J_t(\Sigma_t^t - \Sigma_{t+1}^t)J_t'. \end{aligned} \tag{27}$$

We also define the cross-time quantities:

$$\begin{aligned} M_1 &= \sum_{t=2}^{T} \Sigma_t^T + \mu_t^T \mu_t^{T'}, \\ M_2 &= \sum_{t=1}^{T-1} \Sigma_t^T + \mu_t^T \mu_t^{T'}, \\ M_3 &= \sum_{t=2}^{T} \Sigma_{t,t-1}^T + \mu_t^T \mu_{t-1}^{T\ '}, \qquad \text{where} \\ \Sigma_{t,t-1}^T &= \Sigma_t^t J_{t-1}^T + (\Sigma_{t+1,t}^T - \left\langle A \right\rangle \Sigma_t^t) J_{t-1}^T. \end{aligned} \tag{28}$$

This last element represents $\mathrm{cov}_q(\varphi_t, \varphi_{t-1}|Z_1^T)$.

E Appendix E: M-step: Computation of the Distribution of δ

We will derive here the update equation for δ in the case where $\{\alpha_i\}_{i=1}^k$ is different for each $\{A_{i,i}\}_{i=1}^k$ diagonal entry of A. Extension to the situation where δ is a unique parameter related to a unique variance parameter α is straightforward.

Recalling the prior form for δ defined in equation (8) as well as equation (6), the optimal

form for the variational posterior $q^*(\delta|Z_1^T)$ becomes:

$$\log q^*(\delta|Z_1^T)^{(l+1)} = \sum_{i=1}^{k} \log q^*(\delta_i|Z_1^T)^{(l+1)} = \underbrace{\sum_{i=1}^{k} \log p(\delta_i)}_{\text{term 1}} + \underbrace{\sum_{i=1}^{k} \langle \log p(\alpha_i|\delta_i) \rangle_{q(\alpha_i)^l}}_{\text{term 2}} + \text{ct}, \tag{29}$$

where throughout this Supplementary Material, ct will denote the normalization constant for the given density or log-density. Recalling equation (8), we have that

$$\text{term 1} = \sum_{i=1}^{k} (-\kappa_{p_i} - 1) \log \delta_i - \delta_i^{-1} \beta_{p_i} + \text{ct}, \tag{30}$$

and

$$\text{term 2} = \sum_{i=1}^{k} -c_{p_i} \log\, \delta_i - \delta_i^{-1} \left\langle \alpha_i^{-1} \right\rangle + \text{ct}. \tag{31}$$

The variational posterior for each element $\{\delta_i\}_{i=1}^k$ is therefore an inverse-gamma distribution with shape and scale parameters $\{\kappa_i; \beta_i\}$ defined as:

$$\begin{aligned} \kappa_i &= \kappa_{p_i} + c_{p_i}, \\ \beta_i &= \beta_{p_i} + \left\langle \alpha_i^{-1} \right\rangle. \end{aligned} \tag{32}$$

F Appendix F: M-step: Computation of the Distribution of α

We will derive here the update equation for α in the situation where A is diagonal. Extension to the situation where A is full is straightforward.

Recalling the prior form for each $\{\alpha_i\}_{i=1}^k$ defined in equation (8) as well as equation (6),

the optimal form for the variational posterior $q^*(\alpha_i|\delta_i, Z_1^T)$ becomes:

$$\begin{aligned}\log q^*(\alpha|\delta, Z_1^T)^{(l+1)} &= \sum_{i=1}^{k} \log q^*(\alpha_i|\delta_i, Z_1^T)^{(l+1)} \\ &= \underbrace{\sum_{i=1}^{k} \langle \log p(\alpha_i|\delta_i) \rangle_{q(\delta_i)^{(l)}}}_{\text{term 1}} + \underbrace{\sum_{i=1}^{k} \langle \log p(a_i|\alpha_i) \rangle_{q(a_i)^{(l)}}}_{\text{term 2}} + \text{ct},\end{aligned} \tag{33}$$

where all terms not depending on α are put in the constant term ct. Let us consider the two terms separately. Under the prior specified in equation (8), term 1 simply equals

$$\text{term 1} = \sum_{i=1}^{k} (-c_{p_i} - 1) \log \alpha_i - \alpha_i^{-1} \left\langle \delta_i^{-1} \right\rangle + \text{ct}. \tag{34}$$

Recalling equation (7), term 2 can be rewritten as

$$\begin{aligned}\text{term 2} &= -\sum_{i=1}^{k} \frac{1}{2} \log \alpha_i - \frac{\alpha_i^{-1}}{2} \sum_{i=1}^{k} \left\langle (A_{ii} - m_{A_{ii}})^2 \right\rangle + \text{ct} \\ &= -\sum_{i=1}^{k} \frac{1}{2} \log \alpha_i - \frac{\alpha_i^{-1}}{2} \sum_{i=1}^{k} (\left\langle A_{ii}^2 \right\rangle + m_{A_{ii}}^2 - 2 m_{A_{ii}} \left\langle A_{ii} \right\rangle)^2 + \text{ct} \\ &= -\sum_{i=1}^{k} \frac{1}{2} \log \alpha_i - \frac{\alpha_i^{-1}}{2} \sum_{i=1}^{k} (\theta_i + (\psi_i \theta_i)^2 + m_{A_{ii}}^2 - 2 m_{A_{ii}} \psi_i \theta_i) + \text{ct}.\end{aligned} \tag{35}$$

The whole expression for $q^*(\alpha|\delta, Z_1^T)^{(l+1)}$ then becomes

$$\begin{aligned}q^*(\alpha|\delta, Z_1^T)^{(l+1)} &= \sum_{i=1}^{k} \log q^*(\alpha_i|\delta_i, Z_1^T) \\ &= \sum_{i=1}^{k} (-c_{p_i} - 1) \log \alpha_i - \alpha_i^{-1} \left\langle \delta_i^{-1} \right\rangle - \frac{1}{2} \log \alpha_i \\ &\quad - \alpha_i^{-1} \frac{(\theta_i + (\psi_i \theta_i)^2 + m_{A_{ii}}^2 - 2 m_{A_{ii}} \psi_i \theta_i)}{2} \Big] + \text{ct},\end{aligned} \tag{36}$$

and so the variational posterior for each element $\{\alpha_i\}_{i=1}^k$ is an inverse-gamma distribution with respectively $\{c_i; b_i\}$ shape and scale parameters defined as:

$$\begin{aligned} c_i &= c_{p_i} + \frac{1}{2}, \\ b_i &= \left\langle \delta_i^{-1} \right\rangle + \frac{(\theta_i + (\psi_i \theta_i)^2 + m_{A_{ii}}^2 - 2 m_{A_{ii}} \psi_i \theta_i)}{2}. \end{aligned} \tag{37}$$

G Appendix G: M-step: Computation of the Distribution of A

The optimal form for the variational posterior $q^*(A|\alpha, \delta, Z_1^T)$ is:

$$\begin{aligned} \log q^*(A|\alpha, \delta, Z_1^T)^{(l+1)} = & \underbrace{\langle \log p(A|\alpha) \rangle_{q(\varphi_1^T)^{(l)}, q(Q)^{(l)}, q(\alpha)^{(l)}}}_{\text{term 1}} \\ & + \underbrace{\sum_{t=2}^{T} \langle \log p(\varphi_t|\varphi_{t-1}, A, Q) \rangle_{q(\varphi_1^T)^{(l)}, q(Q)^{(l)}, q(\alpha)^{(l)}}}_{\text{term 2}} + \text{ct}, \end{aligned} \tag{38}$$

where all terms not depending on A are stacked in the constant term ct and the conditional dependence on the data Z_1^T is dropped for sake of brevity. Let us consider the two terms separately. The development of term 2 is the same whatever the form for A (full, diagonal

or proportional to identity). It becomes

$$
\begin{aligned}
\text{term 2} &= \sum_{t=2}^{T} \langle \log p(\varphi_t|\varphi_{t-1}, Q, \alpha)\rangle_{q(\varphi_1^T)^{(l)}, q(Q)^{(l)}} + \text{ct} \\
&= \sum_{t=2}^{T} \left\langle \underbrace{\frac{1}{2}\log|Q|}_{\longrightarrow \text{ct}} - \frac{1}{2}\text{Tr}[(\varphi_t - A\varphi_{t-1})(\varphi_t - A\varphi_{t-1})^{'}Q] \right\rangle_{q(\varphi_1^T)^{(l)}, q(Q)^{(l)}} + \text{ct} \\
&= \sum_{t=2}^{T} \left\langle -\frac{1}{2}\text{Tr}\big[(\varphi_t\varphi_t^{'} - 2A\varphi_{t-1}\varphi_t^{'} + A\varphi_{t-1}\varphi_{t-1}^{'}A^{'})Q\big] \right\rangle_{q(\varphi_1^T)^{(l)}, q(Q)^{(l)}} + \text{ct} \\
&= -\frac{1}{2}\text{Tr}\left\langle [\underbrace{\sum_{t=2}^{T}\left\langle \varphi_t\varphi_t^{'}\right\rangle}_{\longrightarrow \text{ct}} - 2A\underbrace{\sum_{t=2}^{T}\left\langle \varphi_{t-1}\varphi_t^{'}\right\rangle}_{=M_3(28)} + A\underbrace{\sum_{t=2}^{T}\left\langle \varphi_{t-1}\varphi_{t-1}^{'}\right\rangle}_{=M_2(28)} A^{'})Q] \right\rangle_{q(Q)^{(l)}} + \text{ct.}
\end{aligned}
\tag{39}
$$

We can now rewrite (39) in term 2 of equation (38):

$$
\begin{aligned}
\text{term 2} &= -\frac{1}{2}\left\langle \text{Tr}\big[-2AM_3Q + AM_2A^{'}Q\big]\right\rangle_{q(Q)^{(l)}} + \text{ct} \\
&= -\frac{1}{2}\text{Tr}\big[-2AM_3\langle Q\rangle + AM_2A^{'}\langle Q\rangle\big] + \text{ct.}
\end{aligned}
\tag{40}
$$

If we suppose the matrix A diagonal as implied by the prior form in equation (7) as well as equation (6), term 1 in equation (38) is straightforward:

$$
\text{term 1} = \sum_{i=1}^{k} \langle \log p(A_{ii}|\alpha_i)\rangle_{q(\alpha_i)^{(l)}} \quad = -\frac{\left\langle \alpha_i^{-1}\right\rangle}{2}\sum_{i=1}^{k}(A_{ii} - m_{A_{ii}})^2 + \text{ct.} \tag{41}
$$

The full expression then gives:

$$\begin{aligned}
\log q^*(A|\alpha,\delta,Z_1^T)^{(l+1)} &= -\frac{1}{2}\big[\sum_{i=1}^{k}\left\langle\alpha_i^{-1}\right\rangle(A_{ii}-m_{A_{ii}})^2 \\
&\quad + \mathrm{Tr}\{-AM_3\langle Q\rangle - M_3A'\langle Q\rangle + AM_2A'\langle Q\rangle\}\big] + \mathrm{ct} \\
&= -\frac{1}{2}\sum_{i=1}^{k}\big[\left\langle\alpha_i^{-1}\right\rangle(A_{ii}^2+m_{A_{ii}}^2-2A_{ii}m_{A_{ii}}) - 2A_{ii}M_{3_{i,i}}\langle Q_{ii}\rangle \\
&\quad + A_{ii}M_{2_{i,i}}A_{ii}\langle Q_{ii}\rangle\big] + \mathrm{ct} \\
&= -\frac{1}{2}\sum_{i=1}^{k}\big[-2A_{ii}\underbrace{(\left\langle\alpha_i^{-1}\right\rangle m_{A_{ii}} + M_{3_{i,i}}\langle q_i\rangle)}_{:=\psi_i} \\
&\quad + A_{ii}^2\underbrace{(\langle Q_{ii}\rangle M_{2_{i,i}} + \left\langle\alpha_i^{-1}\right\rangle)}_{:={\theta_i}^{-1}} + \underbrace{\left\langle\alpha_i^{-1}\right\rangle m_{A_{ii}}^2}_{\longrightarrow\mathrm{ct}}\big] + \mathrm{ct} \\
&= -\frac{1}{2}\sum_{i=1}^{k}\big[-2A_{ii}\psi_i\theta_i{\theta_i}^{-1} + A_{ii}^2{\theta_i}^{-1}\big] + \mathrm{ct}.
\end{aligned} \tag{42}$$

Therefore

$$\begin{aligned}
&\log q^*(A|\alpha,\delta,Z_1^T) = \sum_{i=1}^{k}\log q^*(A_{ii}|\alpha,\delta,Z_1^T), \\
&\text{with} \qquad q^*(A_{ii}|Z_1^T) = \mathcal{N}_1(A_{ii}|\psi_i\theta_i,\theta_i), \\
&\text{where} \qquad \theta_i = (\langle Q_{ii}\rangle M_{2_{i,i}} + \left\langle\alpha_i^{-1}\right\rangle)^{-1} \qquad \text{and} \\
&\qquad\qquad\quad \psi_i = \left\langle\alpha_i^{-1}\right\rangle m_a + M_{3_{i,i}}\langle Q_{ii}\rangle.
\end{aligned} \tag{43}$$

H Appendix H: M-step: Computation of the Distribution of a_q

We will derive here the update equation for a_q in the situation where the parameter Q_{ii} is unique. Extension to the situation where $\{a_{q_i}\}_{i=1}^k$ is related to different $\{Q_{ii}\}_{i=1}^k$ is

straightforward. Recalling the prior form for a_q in equation (9) as well as equation (6), the optimal form for the variational posterior $q^*(a_q|Z_1^T)$ becomes:

$$\log q^*(a_q|Z_1^T)^{(l+1)} = \underbrace{\log p(a_q)}_{\text{term 1}} + \underbrace{\langle \log p(Q_{ii}|a_q) \rangle_{q(Q_{ii})^{(l+1)}}}_{\text{term 2}} + \text{ct.} \tag{44}$$

Recalling equation (8), we have that

$$\text{term 1} = (-a_{qp} - 1) \log\, a_q - a_q^{-1} b_{qp} + \text{ct} \tag{45}$$

and

$$\text{term 2} = -n_p \log a_q - a_q^{-1} \left\langle q^{-1} \right\rangle + \text{ct.} \tag{46}$$

The variational posterior $\log q^*(a_q|Z_1^T)$ is therefore an inverse-gamma distribution with shape and scale parameters $\{a_{qq}; b_{qq}\}$

$$\begin{aligned} a_{qq} &= a_{qp} + n_p, \\ b_{qq} &= b_{qp} + \left\langle q^{-1} \right\rangle. \end{aligned} \tag{47}$$

I Appendix I: M-step: Computation of the Distribution of Q

If the matrix Q is proportional to identity, i.e. with only one element Q_{ii}, the optimal variational form for the posterior $q^*(Q|a_q)$ becomes:

$$\log q^*(Q|a_q)^{(l+1)} = \underbrace{\log p(Q_{ii}|a_q)}_{\text{term 1}} + \underbrace{\sum_{t=2}^{T} \langle \log p(\varphi_t|A\varphi_{t-1}) \rangle_{q(A)^{(l)}, q(\varphi_1^T)^{(l)}}}_{\text{term 2}} + \text{ct}, \tag{48}$$

where all terms not depending on Q are included in the constant term. Recalling equation (9), we have that

$$\text{term } 1 = (-n_p - 1)\log Q_{ii} - Q_{ii}^{-1}\left\langle a_q^{-1}\right\rangle + \text{ct}, \tag{49}$$

and

$$\begin{aligned}
\text{term } 2 =& -\frac{T-1}{2}\log|Q| - \frac{1}{2}\sum_{t=2}^{T}\Big\langle \text{Tr}\big[(\varphi_t - A\varphi_{t-1})(\varphi_t - A\varphi_{t-1})'Q^{-1}\big]\Big\rangle_{q(\varphi_t,\varphi_{t-1})^{(l)},q(A)^{(l)}} \\
&+ \text{ct} \\
=& -\frac{T-1}{2}\log|Q| \\
&- \frac{1}{2}\left\langle \text{Tr}\big[\sum_{t=2}^{T}\varphi_t\varphi_t^{'}Q^{-1} - 2\sum_{t=2}^{T}\varphi_t\varphi_{t-1}^{'}A^{'}Q^{-1} + A\sum_{t=2}^{T}\varphi_{t-1}\varphi_{t-1}^{'}A^{'}Q^{-1}\big]\right\rangle_{q(\varphi_t,\varphi_{t-1})^{(l)},q(A)^{(l)}} \\
&+ \text{ct} \\
=& -\frac{T-1}{2}\log|Q| \\
&- \frac{1}{2}\left\langle \text{Tr}\big[\left\langle\sum_{t=2}^{T}\varphi_t\varphi_t^{'}\right\rangle Q^{-1} - 2\left\langle\sum_{t=2}^{T}\varphi_t\varphi_{t-1}^{'}\right\rangle A^{'}Q^{-1} + A\left\langle\sum_{t=2}^{T}\varphi_{t-1}\varphi_{t-1}^{'}\right\rangle A^{'}Q^{-1}\big]\right\rangle_{q(A)^{(l)}} \\
&+ \text{ct} \\
=& -k\frac{T-1}{2}\log q - \frac{1}{2}\sum_{i=1}^{k}\left[M_{1_{i,i}} - 2M_{3_{i,i}}\left\langle a_i\right\rangle + M_{2_{i,i}}\left\langle a_i^2\right\rangle\right]Q_{ii}^{-1} + \text{ct} \\
=& -k\frac{T-1}{2}\log Q_{ii} - \frac{1}{2}\underbrace{\sum_{i=1}^{k}\left[M_{1_{i,i}} - 2M_{3_{i,i}}\psi_i\theta_i + M_{2_{i,i}}[(\psi_i\theta_i)^2 + \theta_i^2]\right]}_{:=\Gamma}Q_{ii}^{-1} + \text{ct}.
\end{aligned} \tag{50}$$

The whole expression for $q^*(Q_{ii}|a_q)$ can the be rewritten as

$$\log q^*(Q_{ii}|a_q) = (-n_p - 1)\log Q_{ii} - Q_{ii}^{-1}\left\langle a_q^{-1}\right\rangle - k\frac{T-1}{2}\log Q_{ii} - Q_{ii}^{-1}\frac{\Gamma}{2} + \text{ct}. \tag{51}$$

The variational posterior distribution $q^*(Q_{ii}|a_q)$ is therefore gamma with shape and scale parameters

$$\begin{aligned} n &= n_p + k\frac{T-1}{2}, \\ d &= \left\langle a_q^{-1} \right\rangle + \frac{\Gamma}{2}, \end{aligned} \tag{52}$$

where $\Gamma = \sum_{i=1}^{k} \left[M_{1_{i,i}} - 2M_{3_{i,i}}\psi_i\theta_i + M_{2_{i,i}}[(\psi_i\theta_i)^2 + \theta_i^2] \right]$ and $\{\theta_i; \psi_i\}$ are defined in equation (43).

J Appendix J: M-step: Computation of the Distribution of a_r

We will derive here the update equation for the parameters $\{a_{r_i}\}_{i=1}^d$. Recalling the prior form for $\{a_{r_i}\}_{i=1}^d$ in equation (10) and equation (6), the optimal form for the variational posterior $q^*(\{a_{r_i}\}_{i=1}^d)$ becomes:

$$\log q^*(\{a_{r_i}\}_{i=1}^d)^{(l+1)} = \underbrace{\log p(\{a_{r_i}\}_{i=1}^d)}_{\text{term 1}} + \underbrace{\left\langle \log p(R|\{a_{r_i}\}_{i=1}^d) \right\rangle_{q(R)^{(l)}}}_{\text{term 2}} + \text{ct.} \tag{53}$$

Recalling equation (10), we have that

$$\text{term 1} = \sum_{i=1}^{d} (-a_{pr} - 1) \log a_{r_i} - \sum_{i=1}^{d} a_{r_i}^{-1} b_{pr} + \text{ct}, \tag{54}$$

and

$$\text{term 2} = \frac{r_p}{2} \log |B_p| - \frac{1}{2}\text{Tr}[B_p \left\langle R^{-1} \right\rangle] + \text{ct}, \qquad \text{where} \quad B_p = 2\nu \, \text{diag}[\frac{1}{a_{r_1}} ... \frac{1}{a_{r_d}}]. \tag{55}$$

We therefore have

$$\begin{aligned}\text{term 2} &= \frac{r_p}{2}\log|2\nu\,\mathrm{diag}[\frac{1}{a_{r_1}}...\frac{1}{a_{r_d}}]| - \frac{1}{2}\mathrm{Tr}\left[2\nu\,\mathrm{diag}[\frac{1}{a_{r_1}}...\frac{1}{a_{r_d}}]\left\langle R^{-1}\right\rangle\right] + \mathrm{ct} \\ &= -\frac{r_p}{2}\sum_{i=1}^{d}\log a_{r_i} - \nu\sum_{i=1}^{d} a_{r_i}^{-1}\left\langle R^{-1}_{\{i,i\}}\right\rangle + \mathrm{ct}.\end{aligned} \tag{56}$$

And then regrouping term 1 and term 2 yield to the inverse-gamma distribution for the variational posterior $\log q^*(a_q|Z_1^T)$ with shape and scale parameters $\{a_{qr}; b_{qr}\}$ defined as

$$\begin{aligned}a_{qr} &= a_{pr} + \frac{r_p}{2}, \\ b_{qr} &= b_{pr} + \nu\left\langle R^{-1}_{\{i,i\}}\right\rangle.\end{aligned} \tag{57}$$

K Appendix K: M-step: Computation of the Distribution of R

Recalling the prior form for R in equation (10) and equation (6), the optimal form for the variational posterior $q^*(R|\{a_{r_i}\}_{i=1}^d, Z_1^T)$ is

$$\log q^*(R|\{a_{r_i}\}_{i=1}^d)^{(l+1)} = \underbrace{\left\langle \log p(R|\{a_{r_i}\}_{i=1}^d)\right\rangle_{q(a_r)}}_{\text{term 1}} + \underbrace{\sum_{t=1}^{T}\left\langle \log p(Z_t|C_t\varphi_t)\right\rangle_{q(\varphi_1^T)^{(l)}}}_{\text{term 2}} + \mathrm{ct}, \tag{58}$$

where all terms not depending on R are included in the constant term ct. Let us consider first the term 1. As the prior defined in (10) is inverse-Wishart, term 1 becomes

$$\begin{aligned}\text{term 1} &= -\frac{r_p + d + 1}{2}\log|R| - \frac{1}{2}\mathrm{Tr}(\langle \mathrm{B_p}\rangle\,\mathrm{R}^{-1}) + \mathrm{ct} \\ &= -\frac{r_p + d + 1}{2}\log|R| - \frac{1}{2}\mathrm{Tr}(2\nu\,\mathrm{diag}[\frac{1}{\langle \mathrm{a_{r_1}}\rangle}...\frac{1}{\langle \mathrm{a_{r_d}}\rangle}]\mathrm{R}^{-1}) + \mathrm{ct},\end{aligned} \tag{59}$$

Let us now consider term 2:

$$\begin{aligned}
\text{term } 2 &= \sum_{t=1}^{T} \langle \log p(Z_t|C_t\varphi_t)\rangle_{q(\varphi_1^T)^{(l)}} + \text{ct} \\
&= -\frac{T}{2}\log|R| - \frac{1}{2}\sum_{t=1}^{T}\mathrm{Tr}\Big[\Big\langle (Z_t - C_t\varphi_t)(Z_t - C_t\varphi_t)^{'}\Big\rangle R^{-1}\Big] + \text{ct} \\
&= -\frac{T}{2}\log|R| - \frac{1}{2}\sum_{t=1}^{T}\mathrm{Tr}\Big[\Big\langle (Z_tZ_t^{'} - 2C_t\varphi_tZ_t^{'} + C_tZ_tZ_t^{'}C_t^{'})\Big\rangle R^{-1}\Big] + \text{ct} \\
&= -\frac{T}{2}\log|R| - \frac{1}{2}\mathrm{Tr}\sum_{t=1}^{T}[(Z_t - C_t\mu_t)(Z_t - C_t\mu_t)^{'} + C_t\Sigma_tC_t^{'}]R^{-1} + \text{ct}.
\end{aligned} \tag{60}$$

Assembling terms 1 and 2 in equations (59) and (60) yields to

$$\begin{aligned}
\log q^*(R|\{a_{r_i}\}_{i=1}^d, Z_1^T)^{(l+1)} = &-\frac{r_p + d + 1}{2}\log|R| - \frac{1}{2}\mathrm{Tr}(\langle \mathrm{B_p}\rangle \mathrm{R}^{-1}) - \frac{\mathrm{T}}{2}\log|\mathrm{R}| \\
&- \frac{1}{2}\mathrm{Tr}\underbrace{\sum_{t=1}^{T}[(Z_t - C_t\mu_t)(Z_t - C_t\mu_t)^{'} + C_t\Sigma_tC_t^{'}]}_{:=B} R^{-1} + \text{ct},
\end{aligned} \tag{61}$$

and then

$$\log q^*(R|\{a_{r_i}\}_{i=1}^d, Z_1^T) = \mathcal{IW}_d(R|r_p + T, \langle B_p\rangle + B). \tag{62}$$

L Appendix L: Multiple trials

The model can be modified to deal with N conditionally independent sequences $\{Z_{1\ (j)}^T\}_{j=1}^N$ which are supposed to have the same hidden state. This reflects the case that arises during an event-related experimental paradigm, where many trials on the same condition are measured.

In Beal (2003) and Cassidy (2002), this extension is solved by first estimating the necessary sufficient statistics for each sequence independently in the E-step, and then by

averaging these statistics to get only one set of sufficient statistics, which is representative of the entire set of independent sequences before performing the M-step. However this approach does not take into account the complex dependence of the variational posterior $q(\varphi_1^T|\{Z_{1\ (j)}^T\}_{j=1}^N)$ on the whole dataset $\{Z_{1\ (j)}^T\}_{j=1}^N$. We can however write and solve the full evidence of all the data if we rewrite

$$\begin{cases} \varphi_{t+1} = A\varphi_t + w_t & w_t \sim \mathcal{N}_k(0, Q) \\ \mathbf{Z}_t = \mathbf{C}_t\varphi_t + v_t & v_t \sim \mathcal{N}_d(0, \mathbf{R}) \end{cases} \quad \text{where} \begin{cases} \varphi_t = & \text{vec}\left[\vartheta_{1(t)}, \vartheta_{2(t)}, .., \vartheta_{p(t)}\right]', \\ \mathbf{Z}_t = & (Y_{t(1)} \dots Y_{t(N)} X_{t(1)} \dots X_{t(N)})' \\ \mathbf{C}_t\varphi_t = & \sum_{j=1}^{p} \vartheta_{j(t)} (Y_{t-j(1)} \dots Y_{t-j(N)} X_{t-j(1)} \dots X_{t-j(N)})', \\ \mathbf{R} = & \text{diag}(R_1 \dots R_N)', \end{cases} \tag{63}$$

where $\mathbf{R}$ is block diagonal of dimensions $d^2N \times d^2N$ with each diagonal element R_j being identically distributed. The state equation remains the same as in model (1). We can observe in model (63) that the time-varying parameter vectors $\{\varphi_t\}_{t=1}^T$ are unique, taking into account the whole dataset $\{Z_{1\ (j)}^T\}_{j=1}^N$. The evidence of the complete model can be rewritten as

$$\begin{aligned} p(\{Z_{1\ (j)}^T\}_{j=1}^N, \varphi_1^T, \Omega_1^b) =& p(Z_{1\ (j)}^T\}_{j=1}^N, \varphi_1^T, \{A_{ii}\}_{i=1}^k \{\alpha_i\}_{i=1}^k, \{\delta_i\}_{i=1}^k, Q_{i,i}, a_q, \mathbf{R}, \{a_{r_i}\}_{i=1}^d) \\ =& \prod_{j=1}^{N} \prod_{i=1}^{T} p(Z_t(j)|C_t(j), \varphi_t, R_j) p(\varphi_1) \prod_{i=2}^{T} p(\varphi_t|A, \varphi_{t-1}, Q) \prod_{i=1}^{k} p(A_{ii}|\alpha_i) \\ & \times \prod_{i=1}^{k} p(\alpha_i|\delta_i) \prod_{i=1}^{k} p(\delta_i) p(Q_{i,i}|a_q) p(a_q) p(R_j|\{a_{r_i}\}_{i=1}^d) \prod_{i=1}^{d} p(a_{r_i}). \end{aligned} \tag{64}$$

All the computations done so far can easily be adapted to this new setting. The variational posterior distributions relative to model (63) are now conditional on the whole dataset $\{Z_{1\ (j)}^T\}_{j=1}^N$. A detailed derivation of the model and related variational posterior

distributions with multiple trials can be founded in Section L.

In the DAG depicted in Figure 1, we can see that conditional dependence on the observed variables Z is present only for nodes $\{\varphi_1^T; Z_1^T; \mathbf{R}; a_r\}$. Figure L.1 therefore shows the part of the DAG in Figure 1 which is modified if multiple sequences are taken into account in the estimation procedure. For the hidden variables, the variational E-step KRTS algorithm can be run on the whole state space system (63) and the resulting variational posterior density will therefore be conditional on the whole dataset $q(\varphi_1^T|\{Z_{1\ (j)}^T\}_{j=1}^N)$. In Appendix D, the update equations remain the same, except for the Kalman gain in equation (26) which becomes

$$K_t = \Sigma_t^{t-1} C_t' (C_t \Sigma_t^{t-1} C_t' + \langle \mathbf{R} \rangle)^{-1}. \tag{65}$$

Looking at the update equations for the a_r parameter in Appendix J, we see that for multiple sequences equation (53) becomes

$$\log q^*(\{a_{r_i}\}_{i=1}^d)^{(l+1)} = \underbrace{\log p(\{a_{r_i}\}_{i=1}^d)}_{\text{term 1}} + \underbrace{\sum_{j=1}^N \left\langle \log p(R_j|\{a_{r_i}\}_{i=1}^d) \right\rangle_{q(R_j)^{(l)}}}_{\text{term 2}} + \text{ct}, \tag{66}$$

where ct denotes the normalization constant for the given density. Therefore equation (56) can be rewritten as

$$\text{term 2} = -\frac{r_p}{2} \sum_{j=1}^N \sum_{i=1}^d \log a_{r_i} - \nu \sum_{j=1}^N \sum_{i=1}^d a_{r_i}^{-1} \left\langle R_{j\{i,i\}}^{-1} \right\rangle + \text{ct}. \tag{67}$$

The resulting shape and scale parameters for $q^*(a_{r_i})^{(l+1)}$ therefore become

$$\begin{aligned} a_{qr} &= a_{pr} + \frac{r_p N}{2}, \\ b_{qr} &= b_{pr} + \nu \sum_{j=1}^N \left\langle R_{j\{i,i\}}^{-1} \right\rangle, \quad \text{where} \quad \left\langle R_{j\{i,i\}}^{-1} \right\rangle = r \left\langle B_j^{-1} \right\rangle. \end{aligned} \tag{68}$$

Looking at the update equations for the R matrix in Appendix K, we see that for multiple sequences, equation (58) becomes

$$\begin{aligned}\log q^*(\mathbf{R}|\{a_{r_i}\}_{i=1}^d)^{(l+1)} &= \sum_{j=1}^{N} \log q^*(R_j|\{a_{r_i}\}_{i=1}^d)^{(l+1)} \\ &= \sum_{j=1}^{N} \underbrace{\left\langle \log p(R_j|\{a_{r_i}\}_{i=1}^d) \right\rangle_{q(a_r)}}_{\text{term 1}} \\ &\quad + \underbrace{\sum_{j=1}^{N}\sum_{t=1}^{T} \langle \log p(Z_t(j)|C_t(j)\varphi_t) \rangle_{q(\varphi_1^T)^{(l)}}}_{\text{term 2}} + \text{ct},\end{aligned} \tag{69}$$

where each block-diagonal entry R_j is conditional on a unique dataset $\{Z_{1\ (j)}^T\}$, and will therefore present the same properties as the unique R matrix in Appendix K. The resulting parameters for the variational posterior inverse-Wishart distribution $q^*(R_j|\{a_{r_i}\}_{i=1}^d)^{(l+1)}$ can therefore be rewritten as

$$\begin{aligned} B_j &= 2\,\nu\,\text{diag}[\frac{1}{\langle a_{r_1}\rangle}, \ldots, \frac{1}{\langle a_{r_d}\rangle}] + \sum_{t=1}^{T}[(Z_t - C_t\mu_t)(Z_t - C_t\mu_t) + C_t\Sigma_t C_t^{'} R], \\ r &= r_p + T. \end{aligned} \tag{70}$$

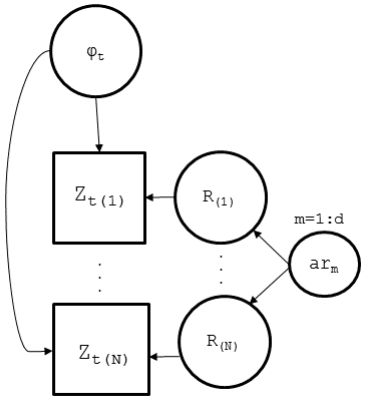

Figure L.1: Part of the DAG representation of the full model given in Figure 1 that is modified to account for multiple trials.

M Appendix M: Granger-Causality Detection Results

M.1 Data generated with slowly-varying parameters and normal errors

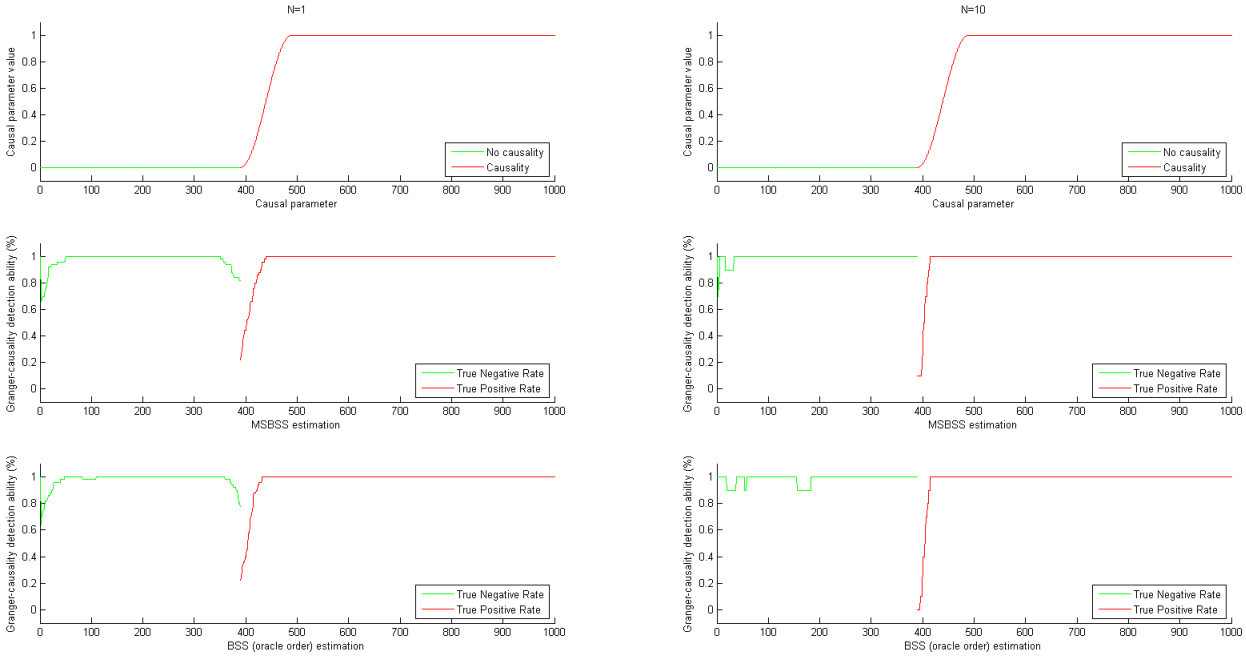

Figure M.1: Granger-causality detection ability for order 1, series length 1000 and causal parameter 1.

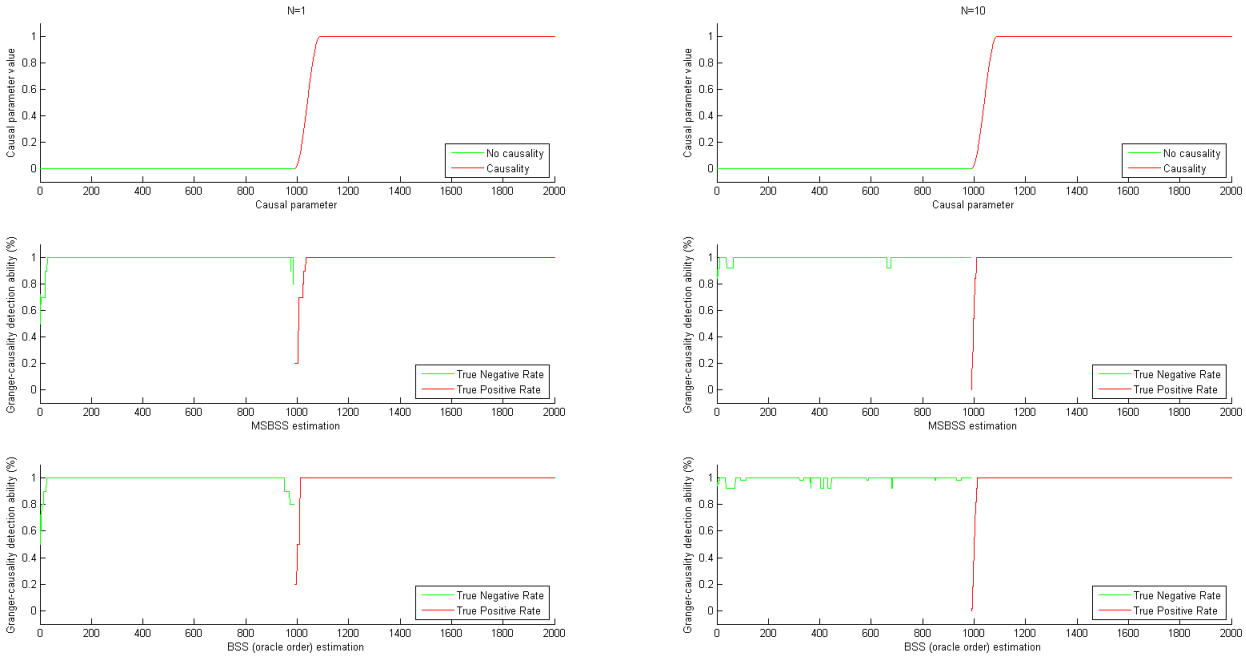

Figure M.2: Granger-causality detection ability for order 1, series length 2000 and causal parameter 1.

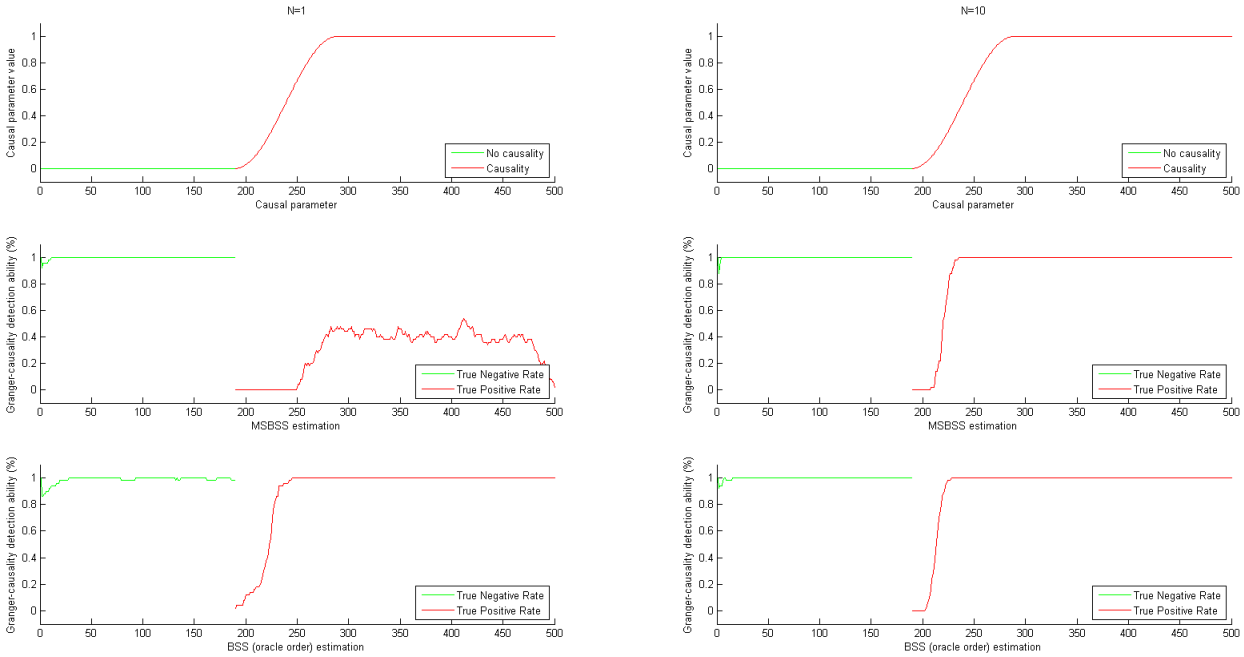

Figure M.3: Granger-causality detection ability for order 2, series length 500 and causal parameter 1.

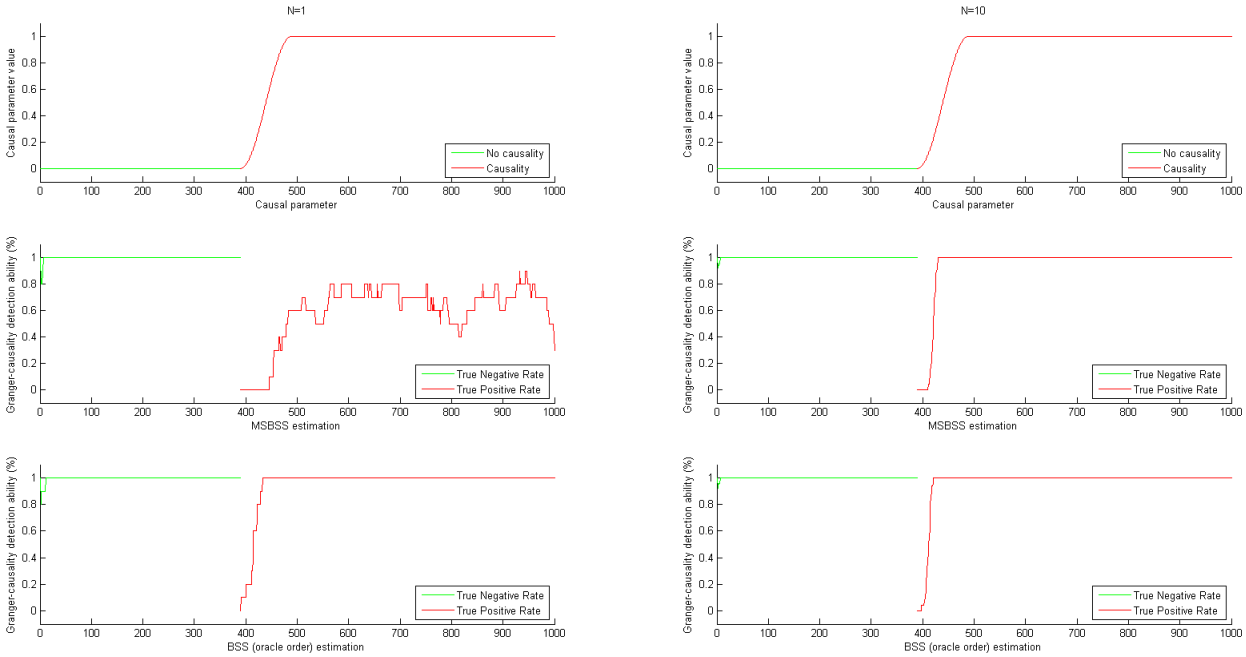

Figure M.4: Granger-causality detection ability for order 2, series length 1000 and causal parameter 1.

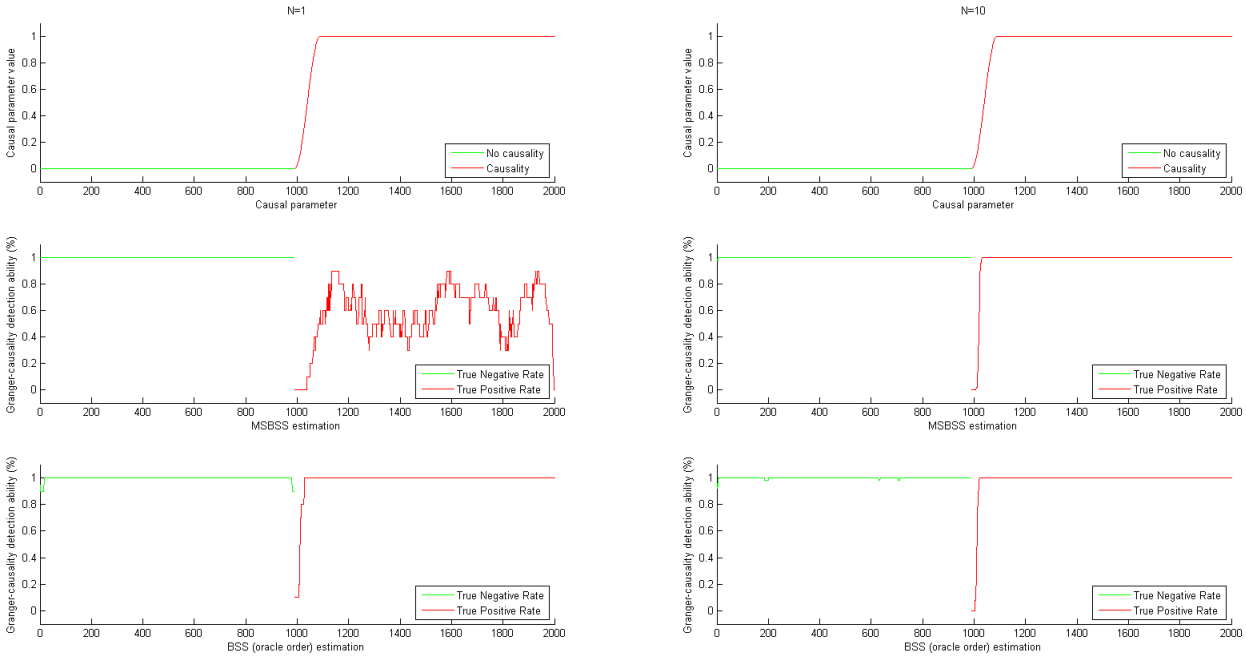

Figure M.5: Granger-causality detection ability for order 2, series length 2000 and causal parameter 1.

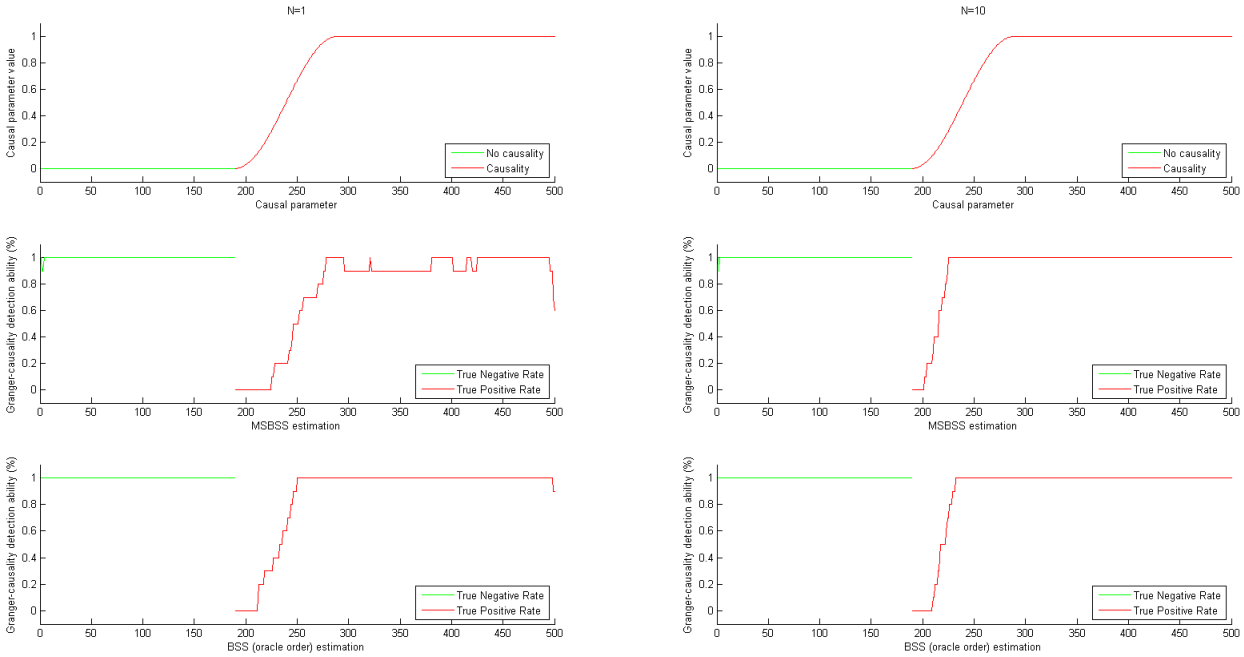

Figure M.6: Granger-causality detection ability for order 4, series length 500 and causal parameter 1.

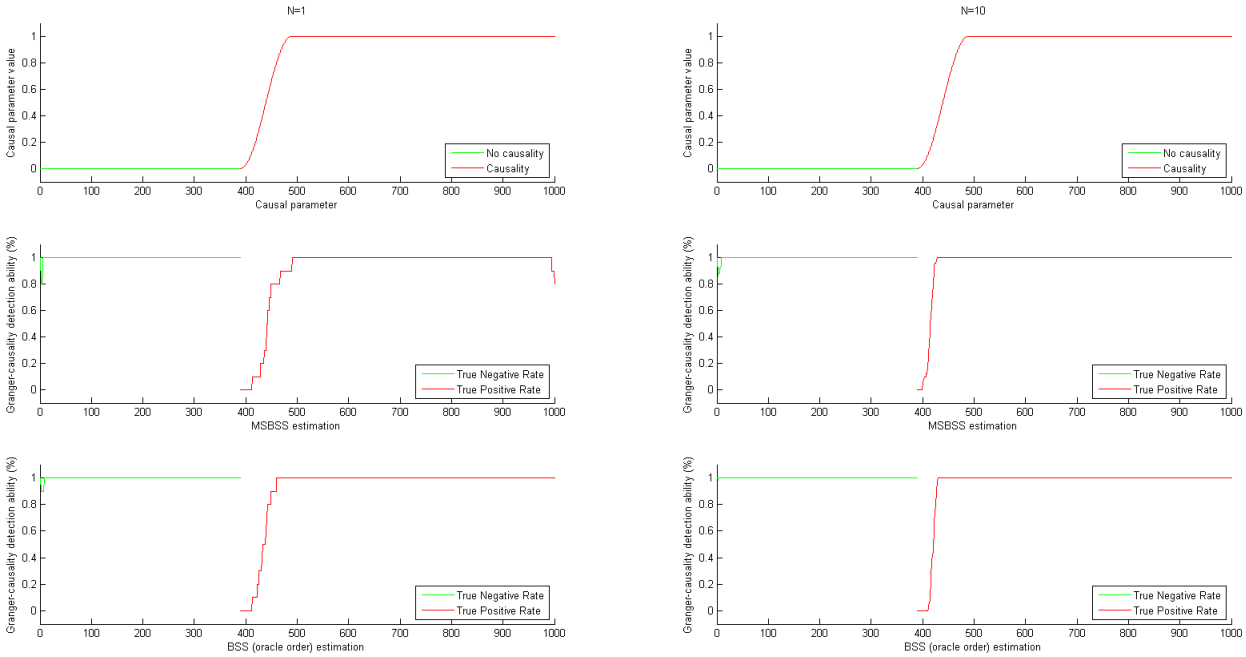

Figure M.7: Granger-causality detection ability for order 4, series length 1000 and causal parameter 1.

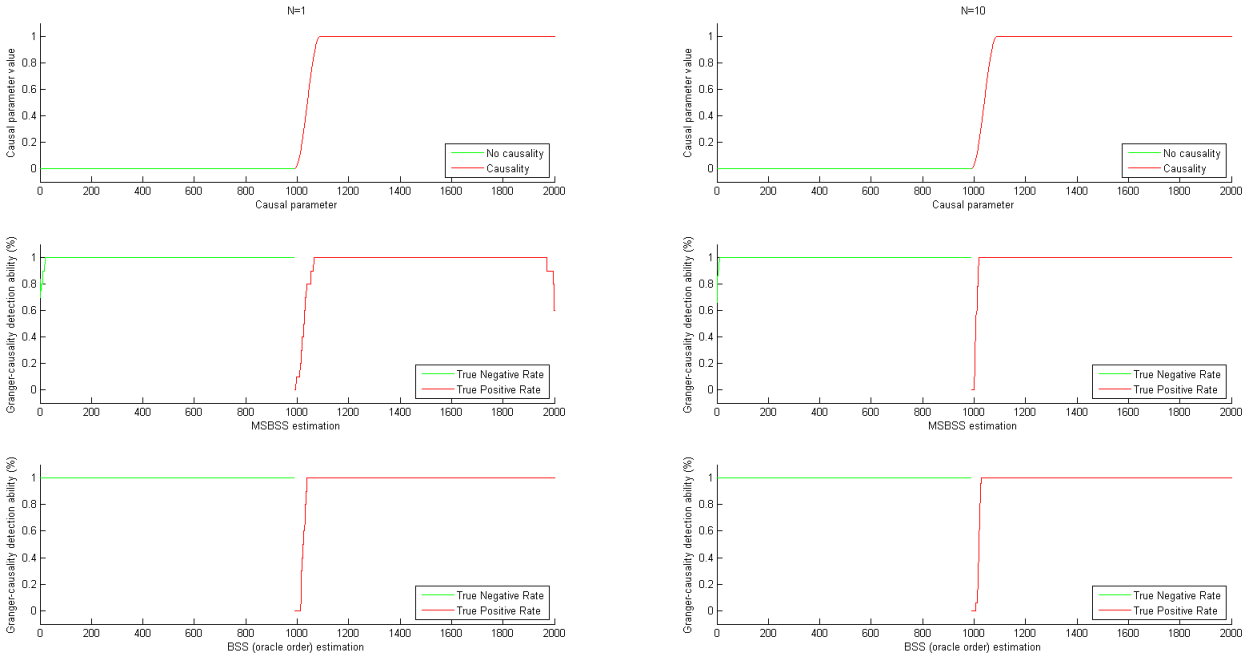

Figure M.8: Granger-causality detection ability for order 4, series length 2000 and causal parameter 1.

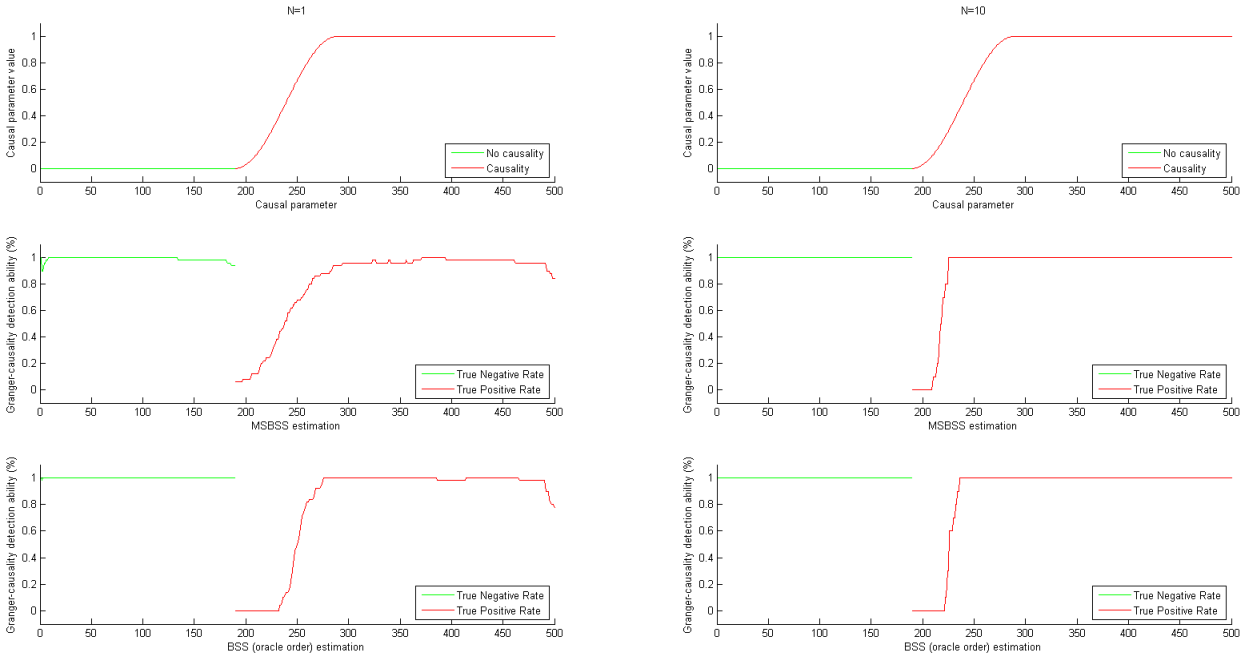

Figure M.9: Granger-causality detection ability for order 8, series length 500 and causal parameter 1.

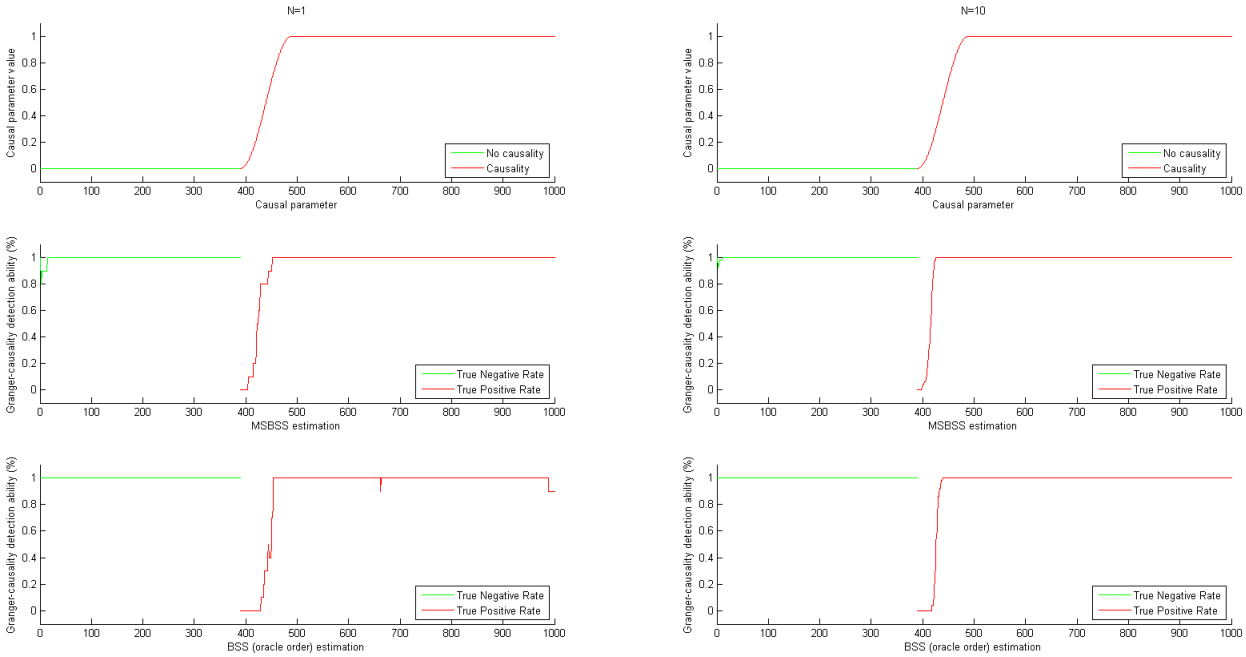

Figure M.10: Granger-causality detection ability for order 8, series length 1000 and causal parameter 1.

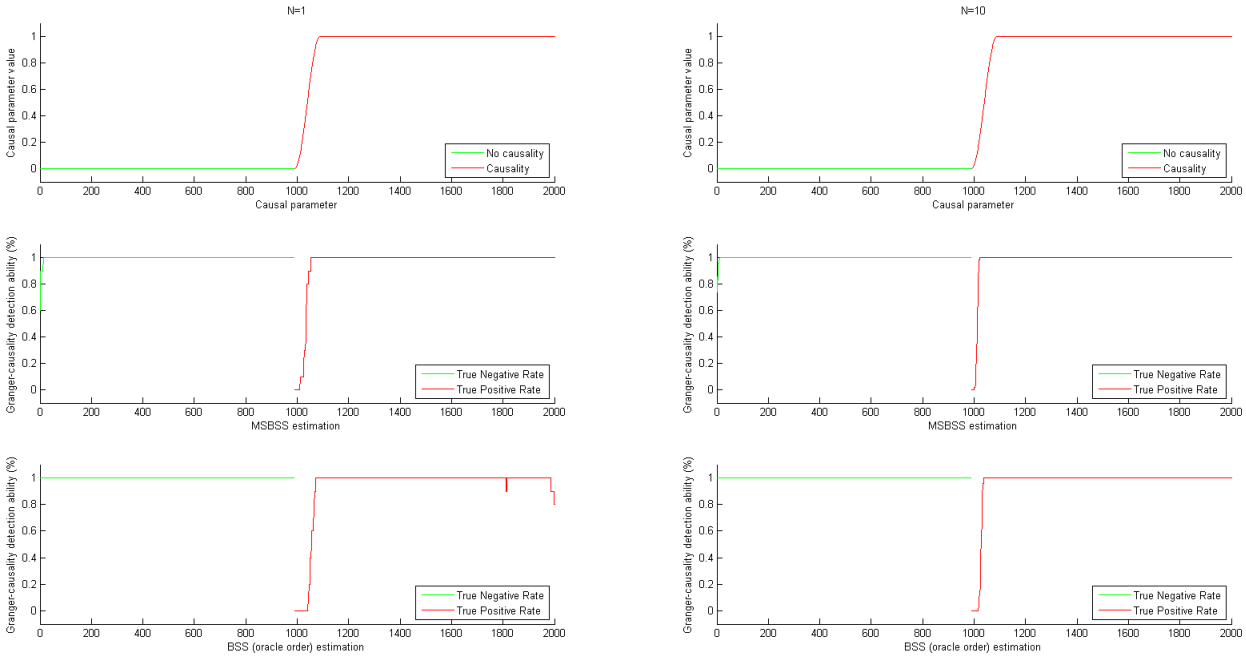

Figure M.11: Granger-causality detection ability for order 8, series length 2000 and causal parameter 1.

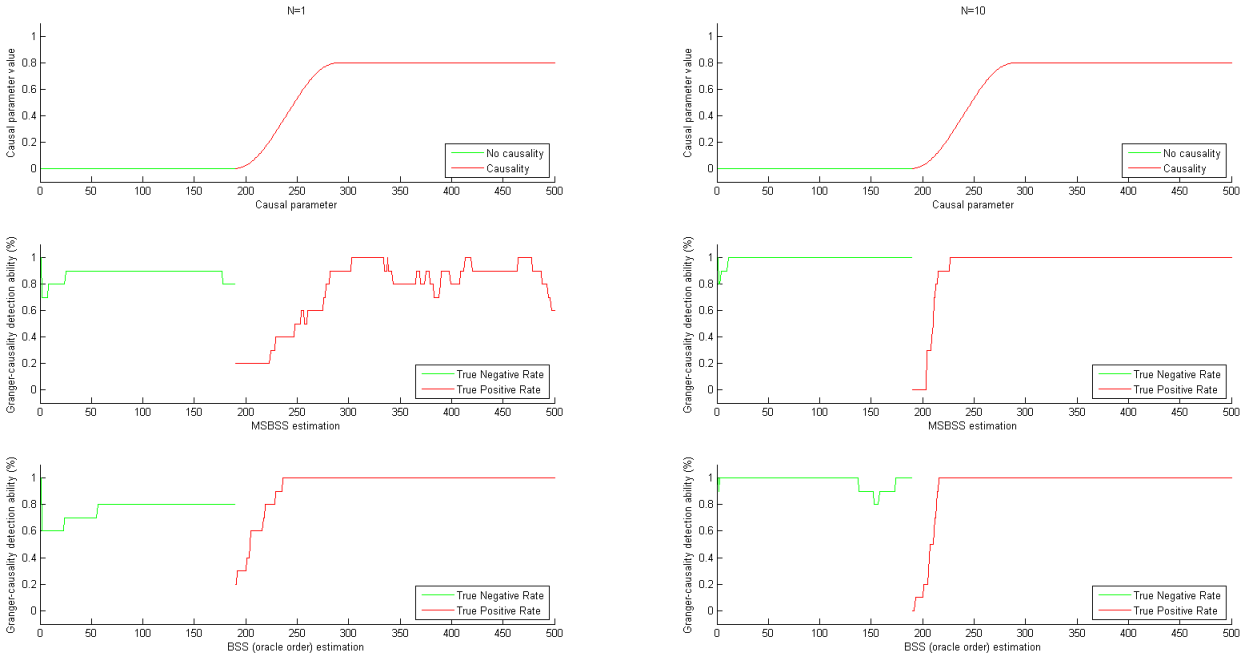

Figure M.12: Granger-causality detection ability for order 1, series length 500 and causal parameter 0.8.

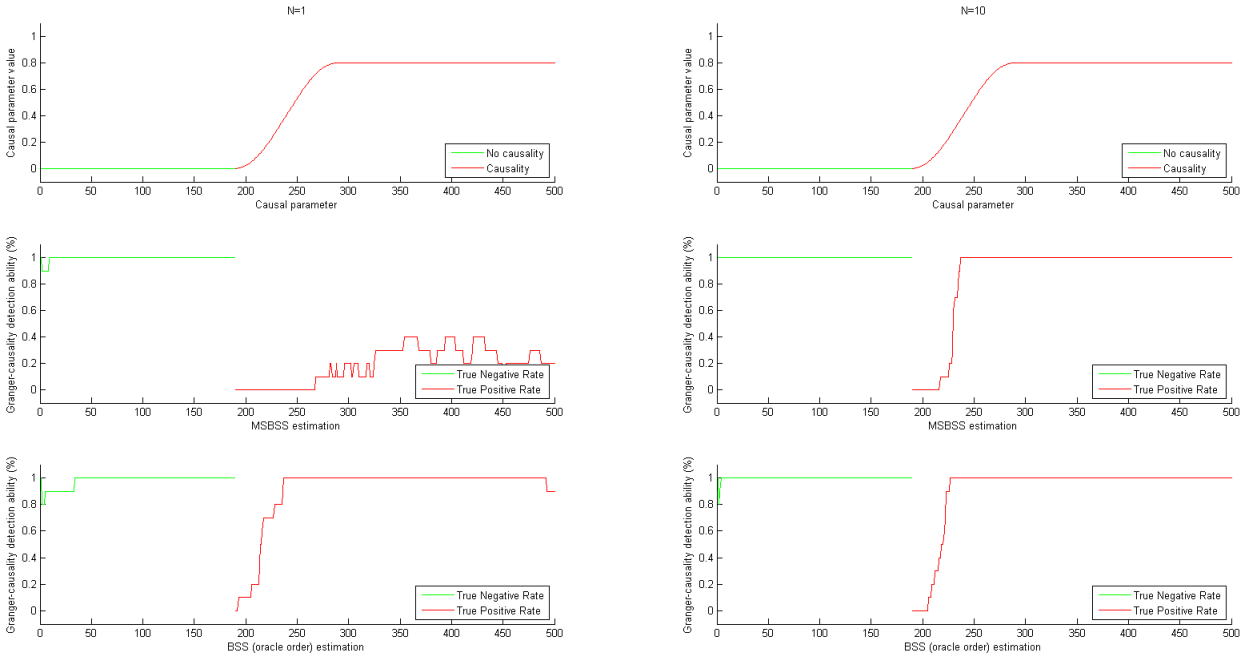

Figure M.13: Granger-causality detection ability for order 2, series length 500 and causal parameter 0.8.

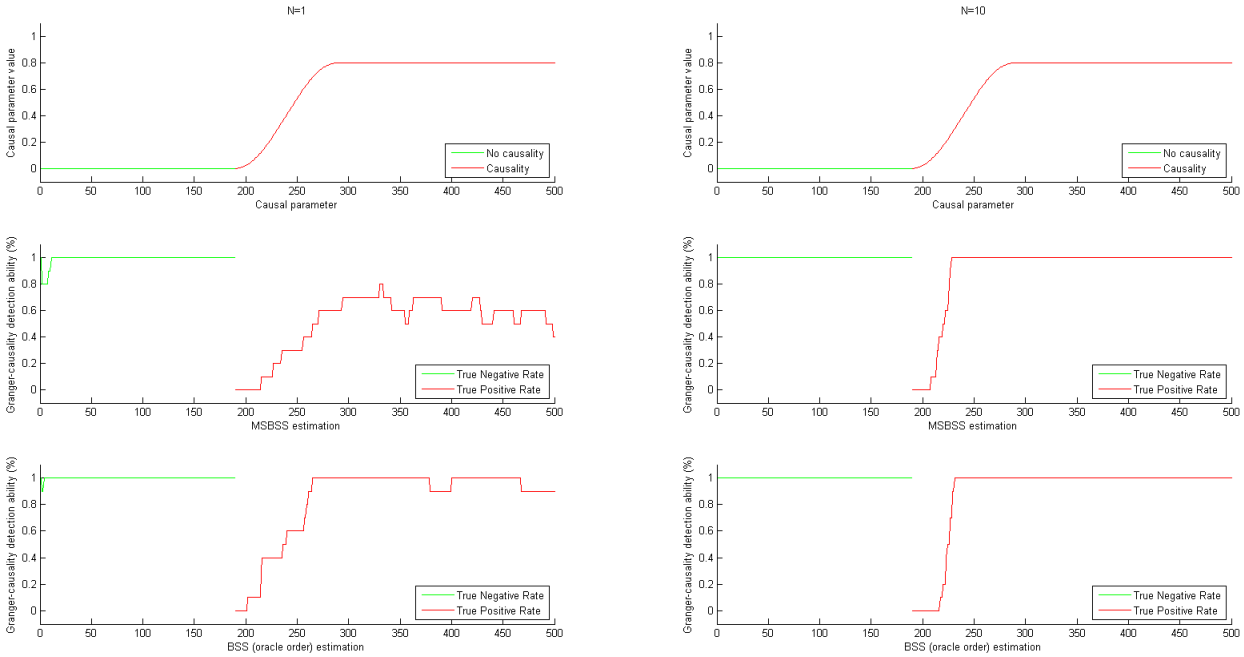

Figure M.14: Granger-causality detection ability for order 4, series length 500 and causal parameter 0.8.

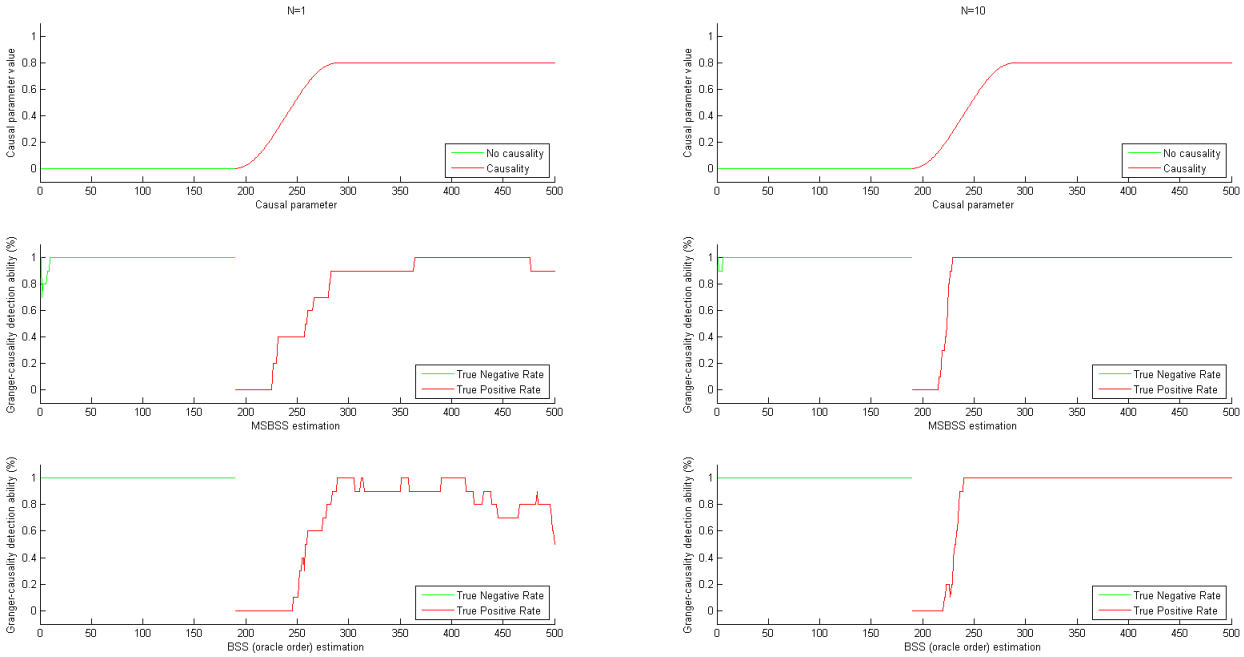

Figure M.15: Granger-causality detection ability for order 8, series length 500 and causal parameter 0.8.

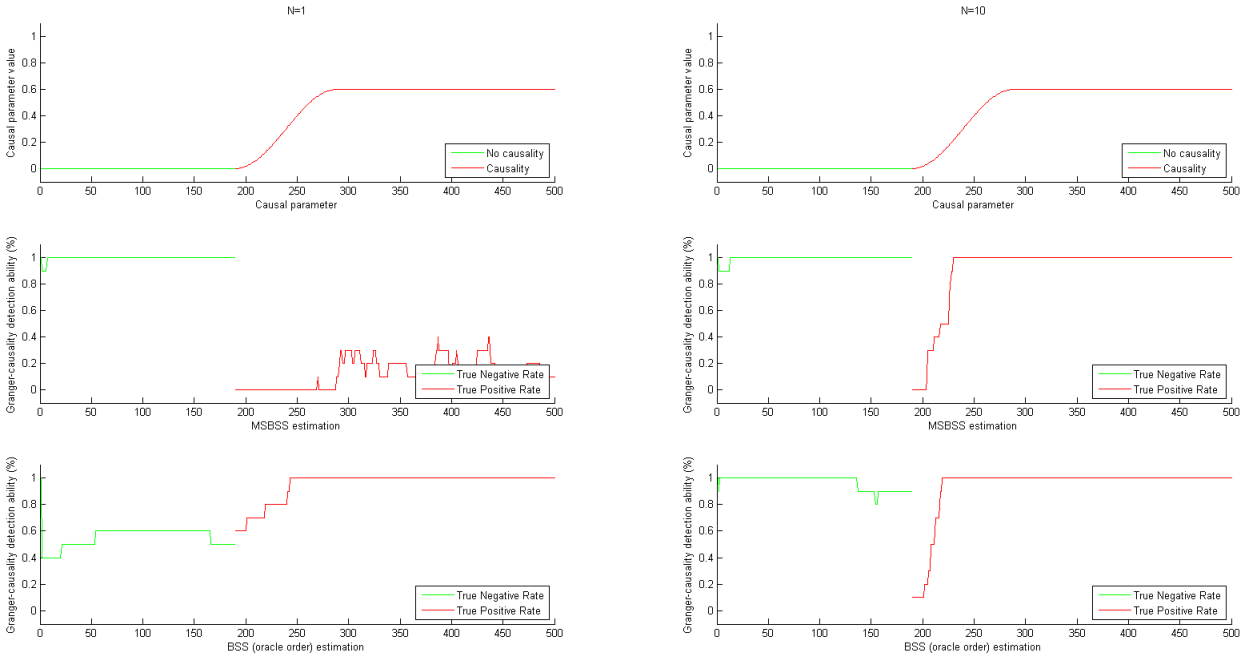

Figure M.16: Granger-causality detection ability for order 1, series length 500 and causal parameter 0.6.

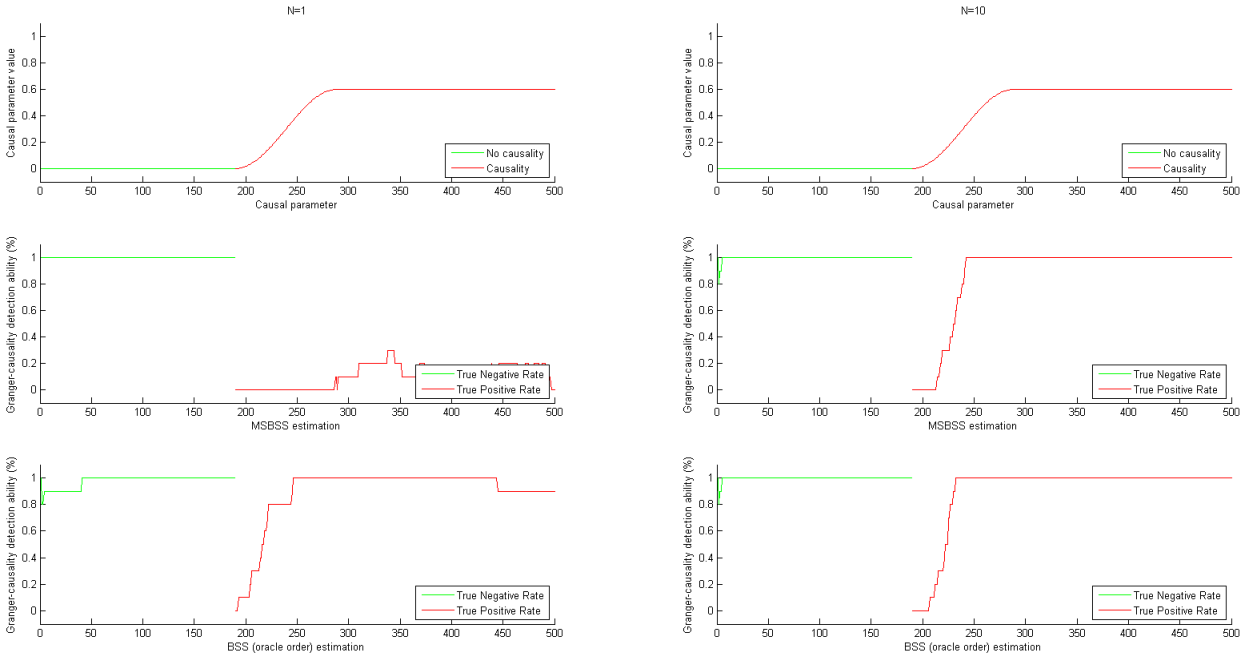

Figure M.17: Granger-causality detection ability for order 2, series length 500 and causal parameter 0.6.

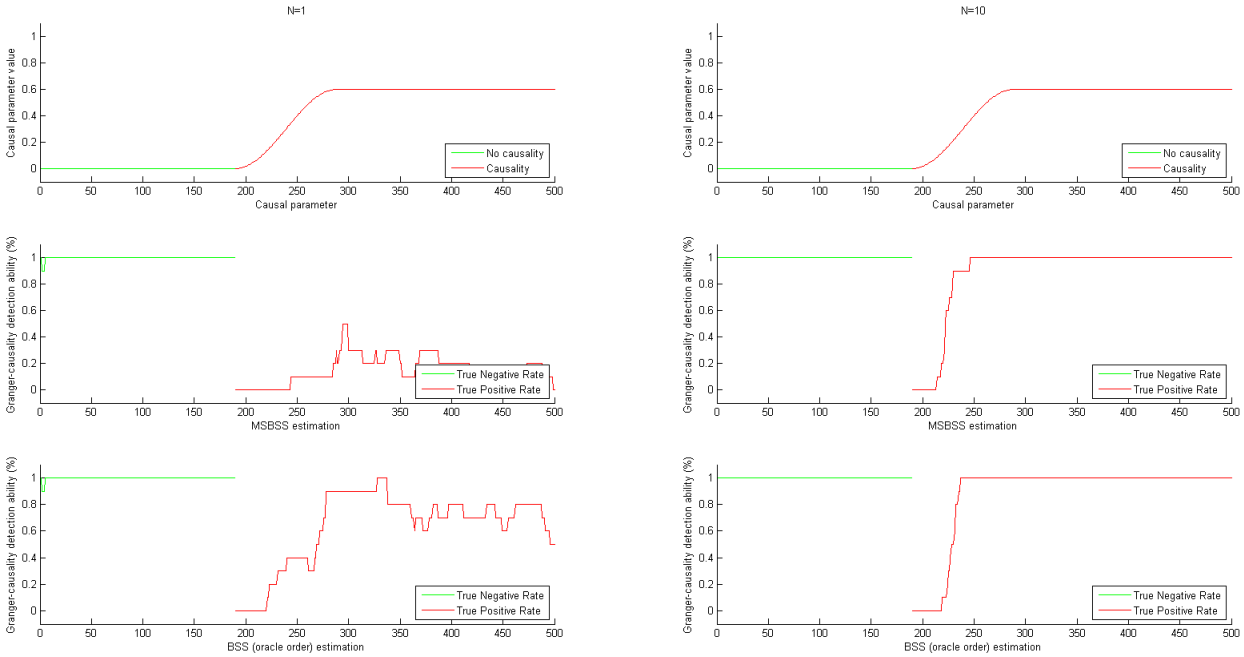

Figure M.18: Granger-causality detection ability for order 4, series length 500 and causal parameter 0.6.

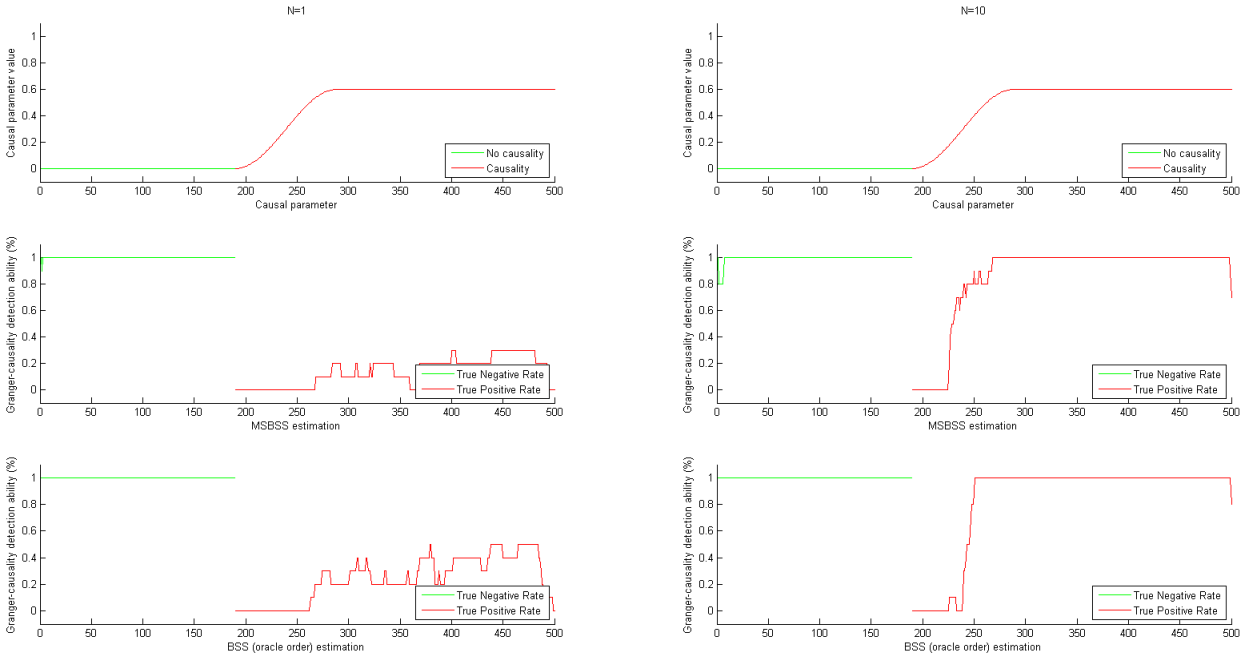

Figure M.19: Granger-causality detection ability for order 8, series length 500 and causal parameter 0.6.

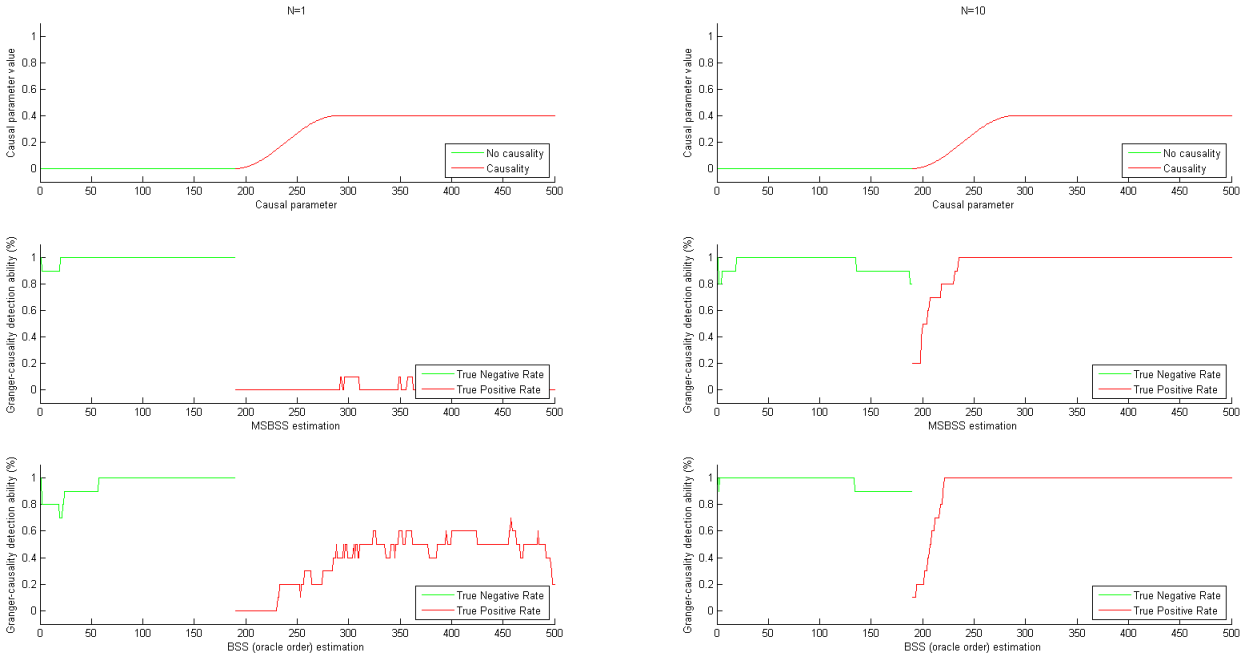

Figure M.20: Granger-causality detection ability for order 1, series length 500 and causal parameter 0.4.

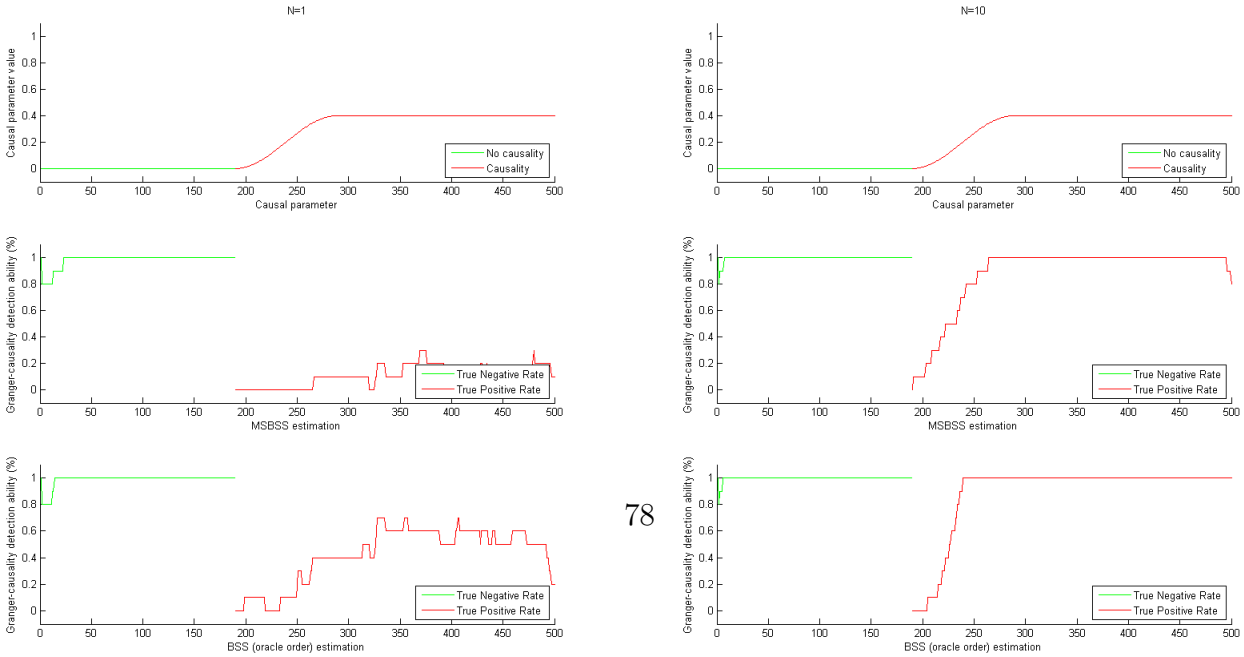

Figure M.21: Granger-causality detection ability for order 2, series length 500 and causal parameter 0.4.

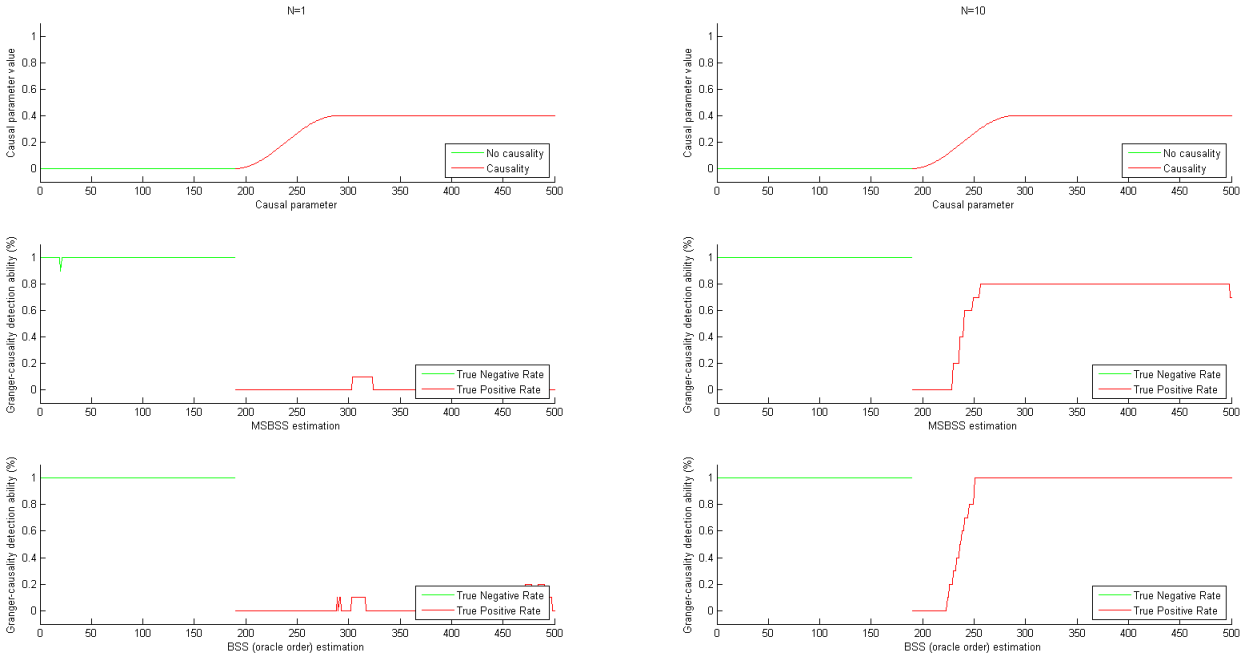

Figure M.22: Granger-causality detection ability for order 4, series length 500 and causal parameter 0.4.

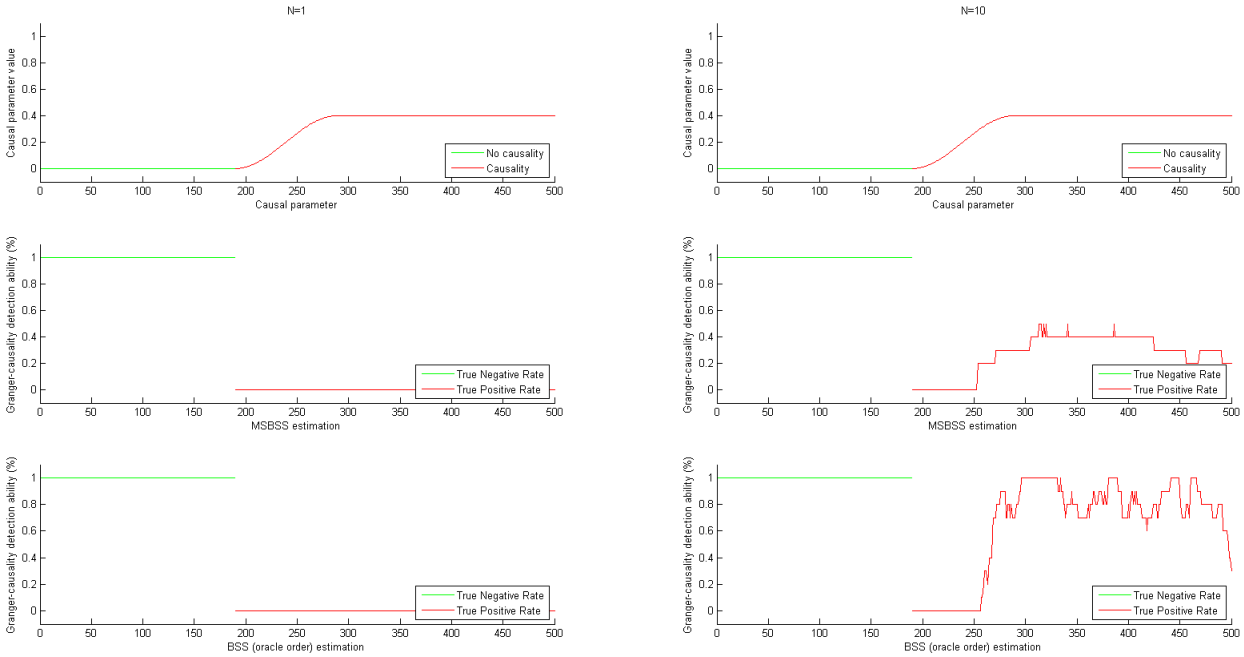

Figure M.23: Granger-causality detection ability for order 8, series length 500 and causal parameter 0.4.

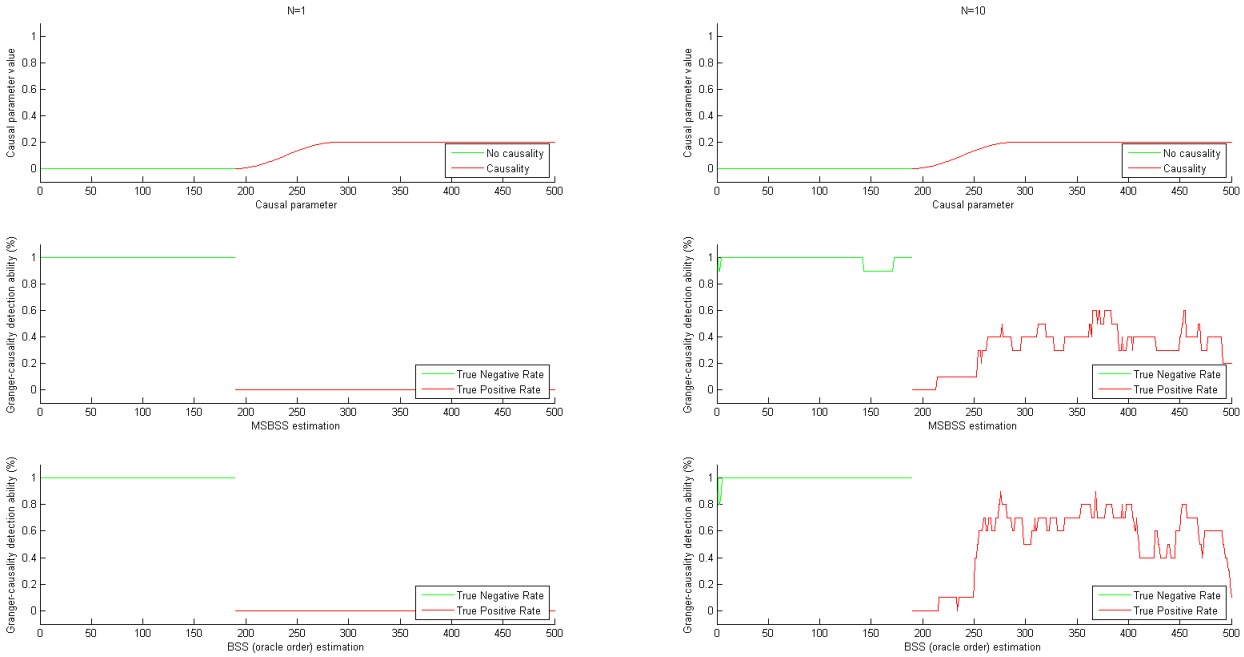

Figure M.24: Granger-causality detection ability for order 2, series length 500 and causal parameter 0.2.

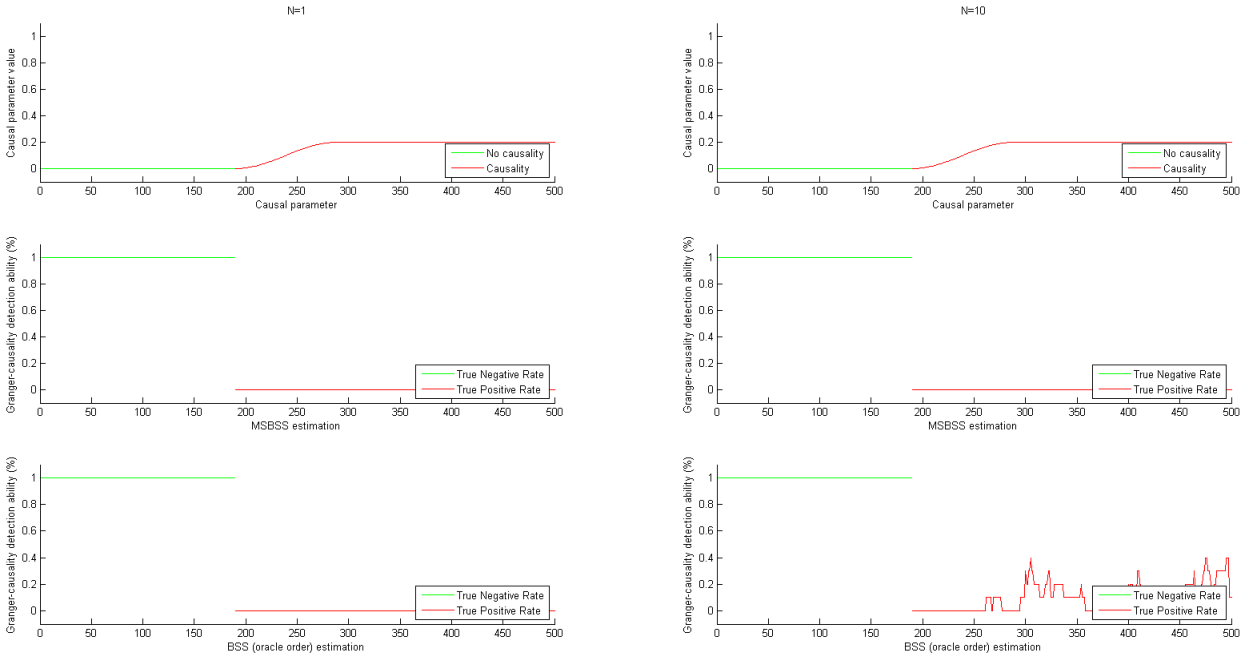

Figure M.25: Granger-causality detection ability for order 4, series length 500 and causal parameter 0.2.

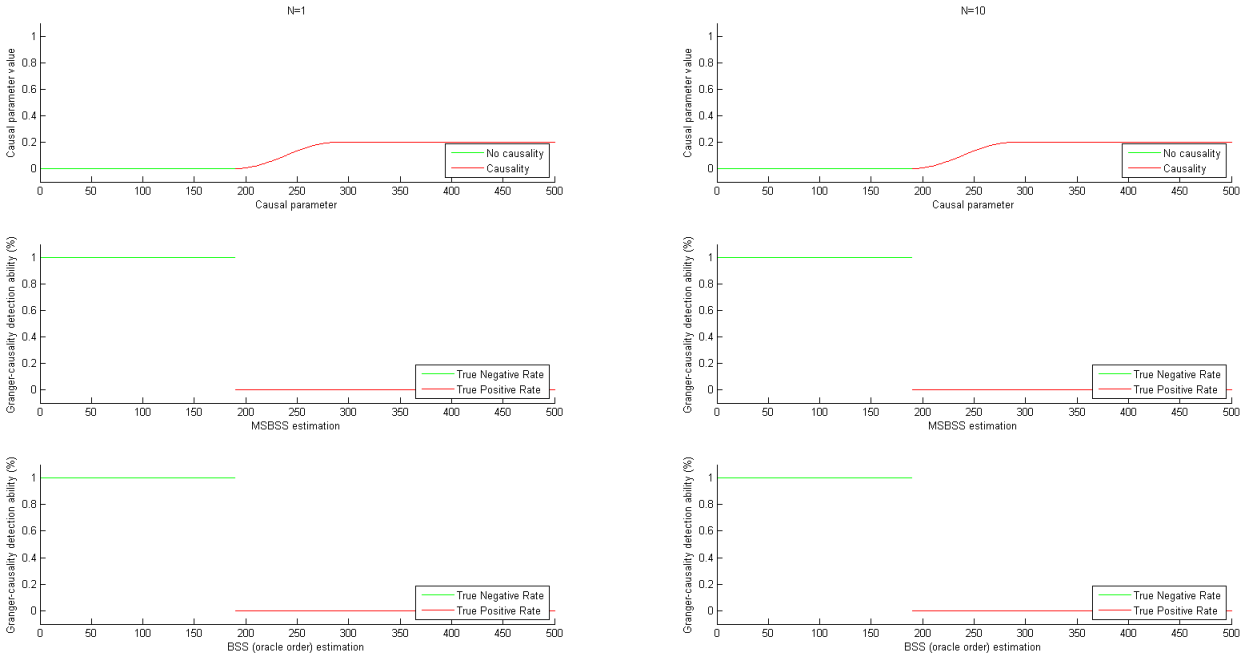

Figure M.26: Granger-causality detection ability for order 8, series length 500 and causal parameter 0.2.

M.2 Data generated with slowly-varying parameters and non-normal errors

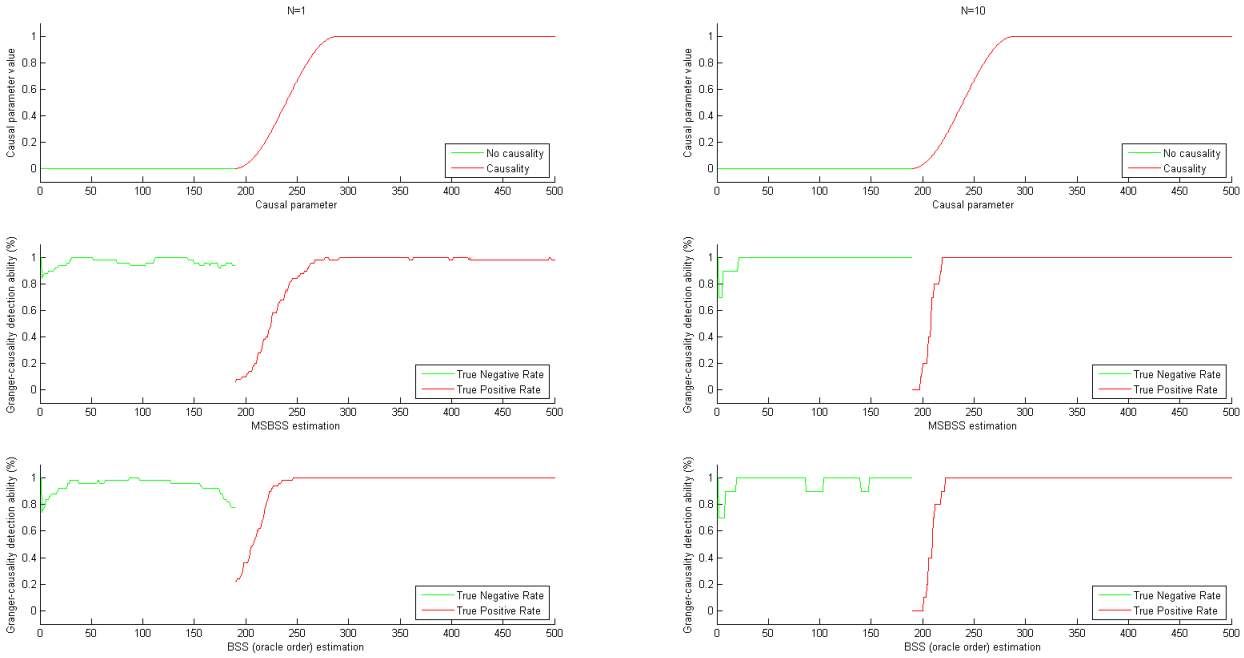

Figure M.27: Granger-causality detection ability for a non-normal error simulation of order 1, series length 500 and causal parameter 1.

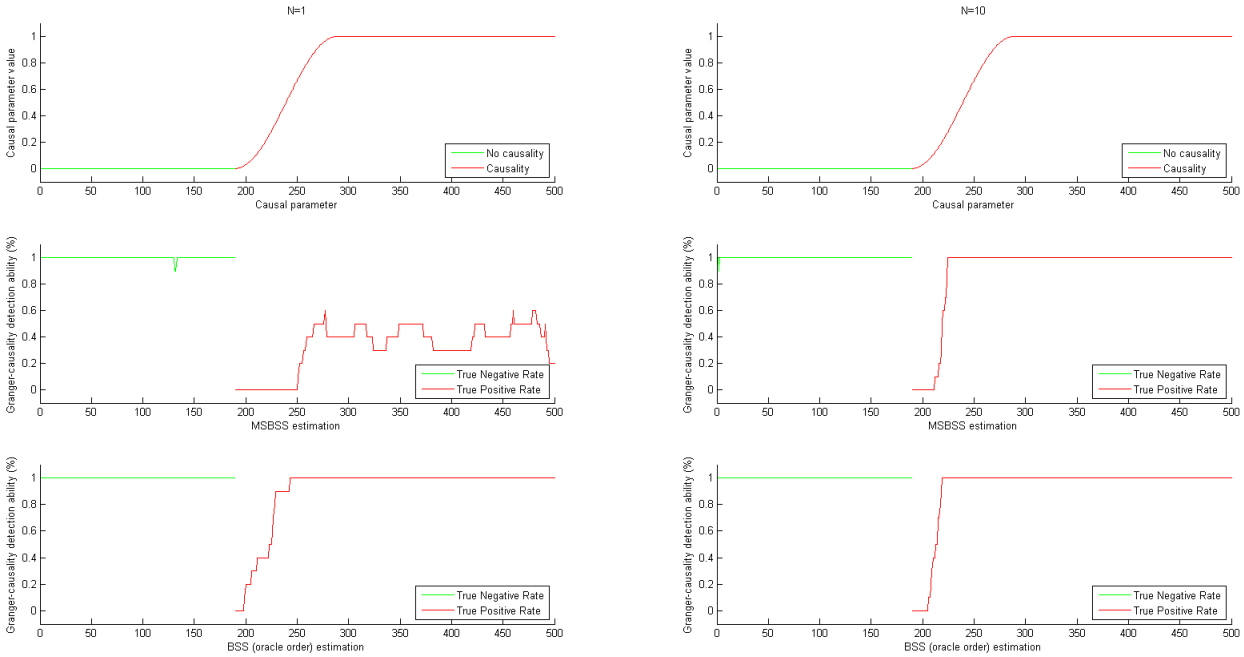

Figure M.28: Granger-causality detection ability for a non-normal error simulation of order 2 and series length 500 and causal parameter 1.

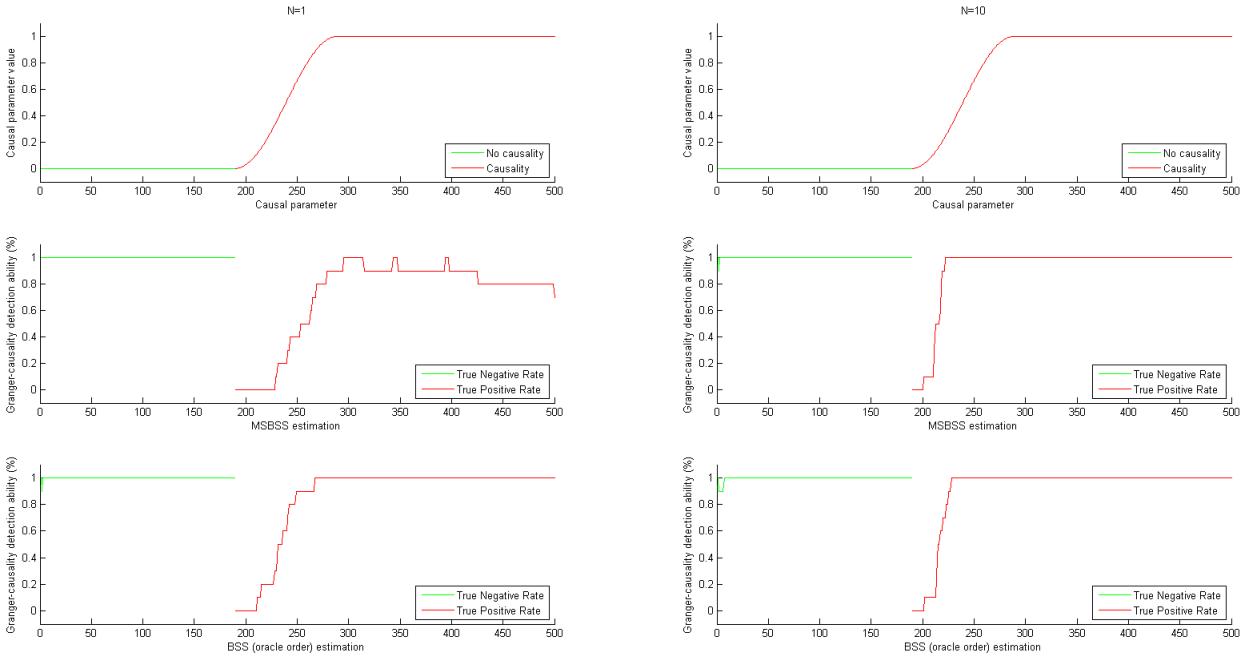

Figure M.29: Granger-causality detection ability for a non-normal error simulation of order 4 and series length 500 and causal parameter 1.

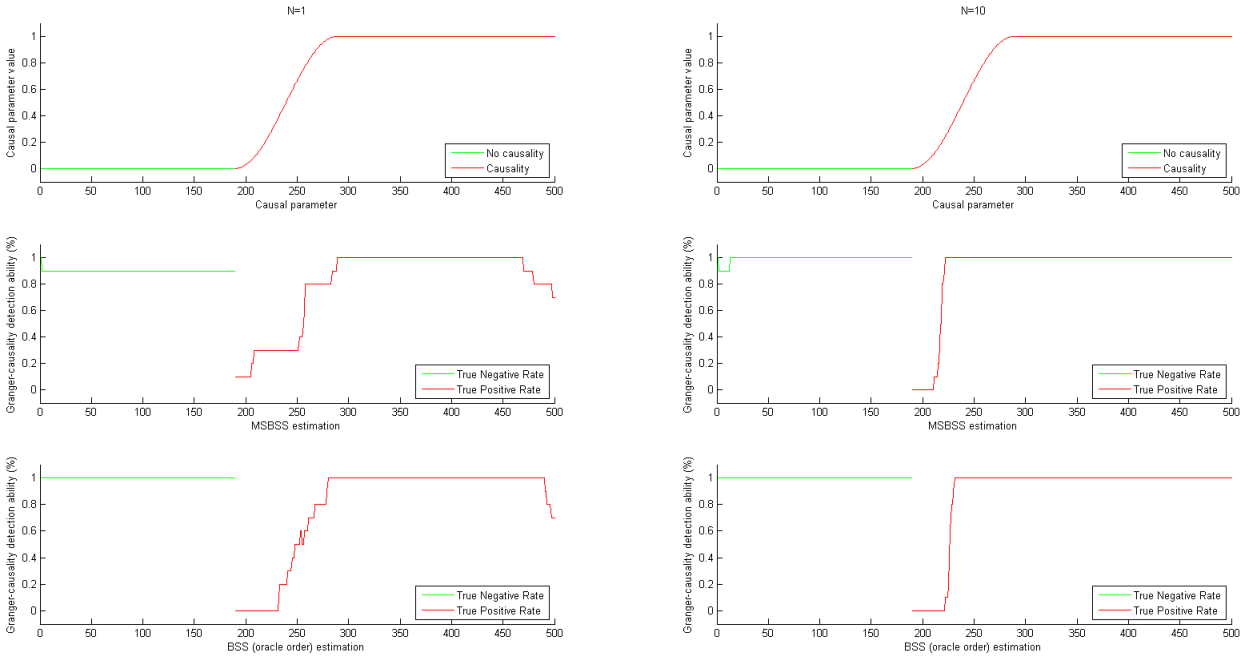

Figure M.30: Granger-causality detection ability for a non-normal error simulation of order 8 and series length 500 and causal parameter 1.

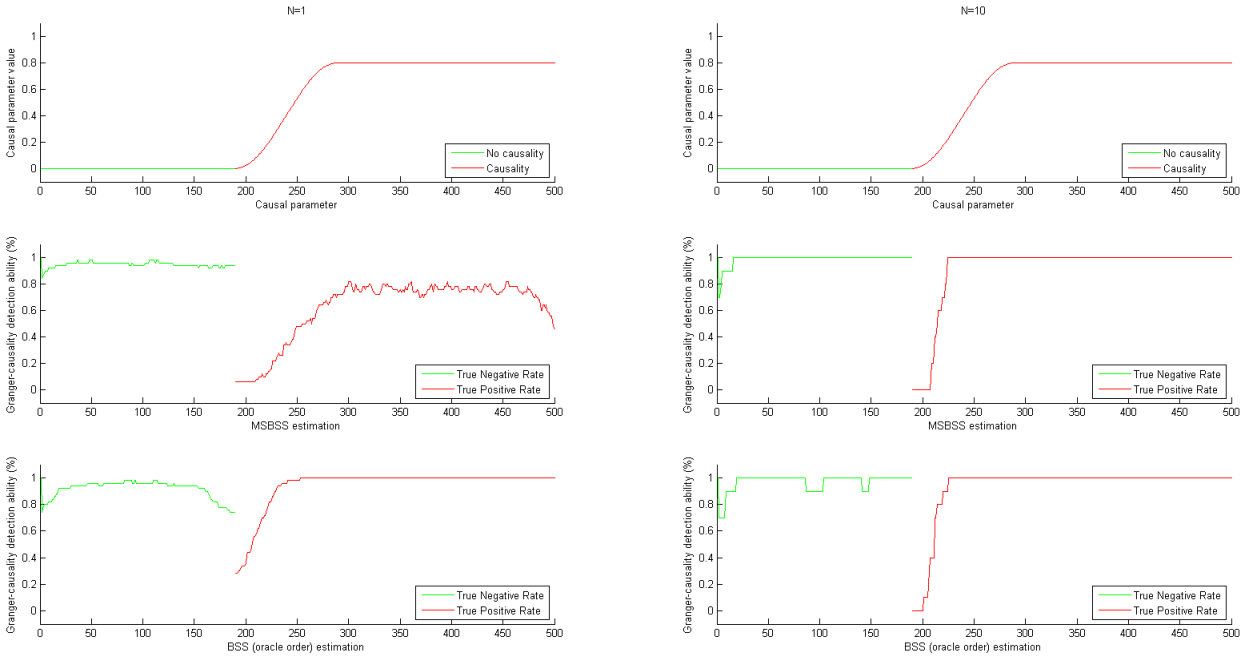

Figure M.31: Granger-causality detection ability for a non-normal error simulation of order 1 and series length 500 and causal parameter 0.8.

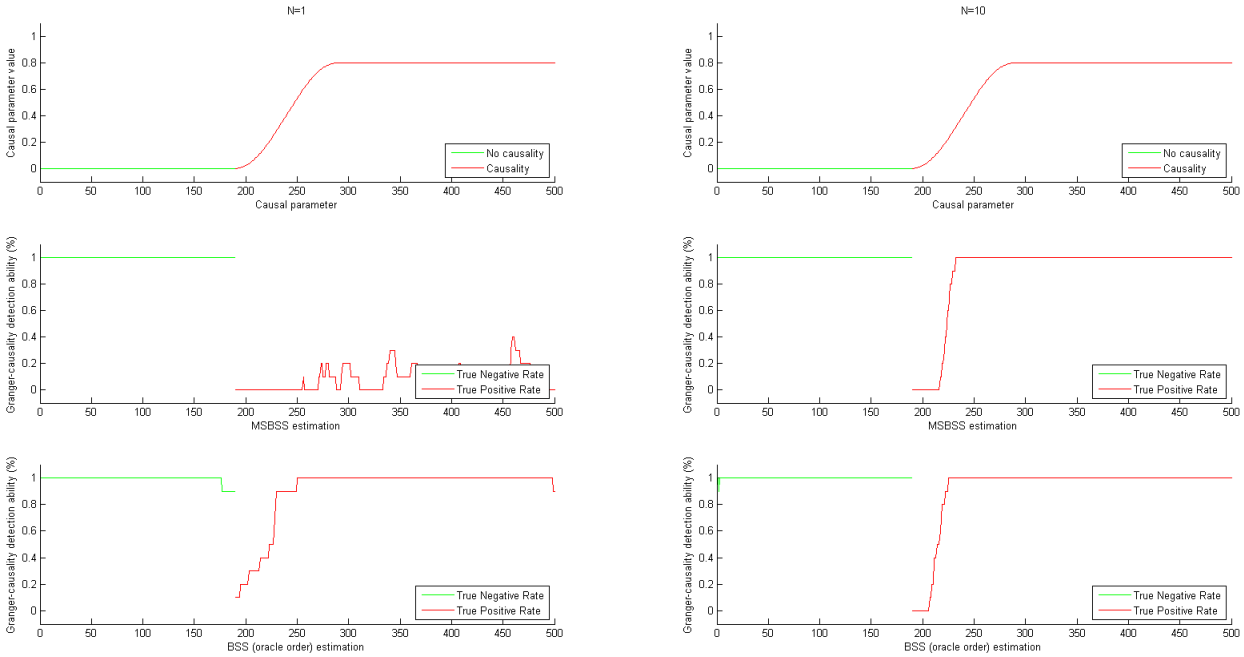

Figure M.32: Granger-causality detection ability for a non-normal error simulation of order 2 and series length 500 and causal parameter 0.8.

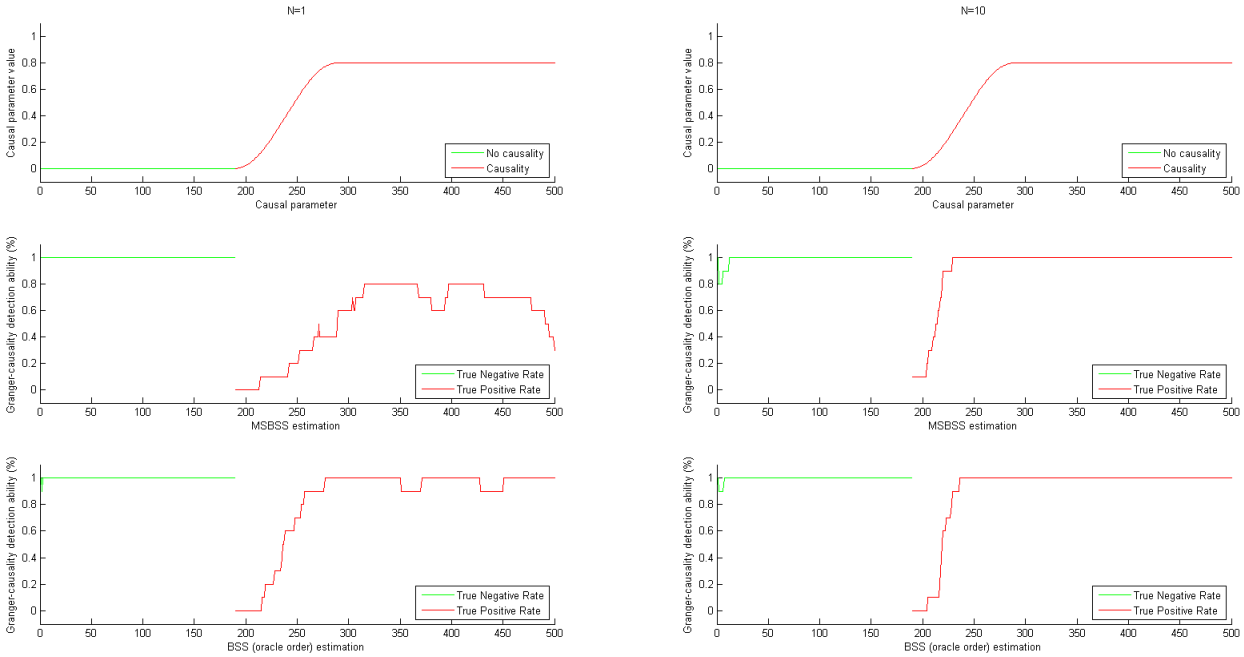

Figure M.33: Granger-causality detection ability for a non-normal error simulation of order 4 and series length 500 and causal parameter 0.8.

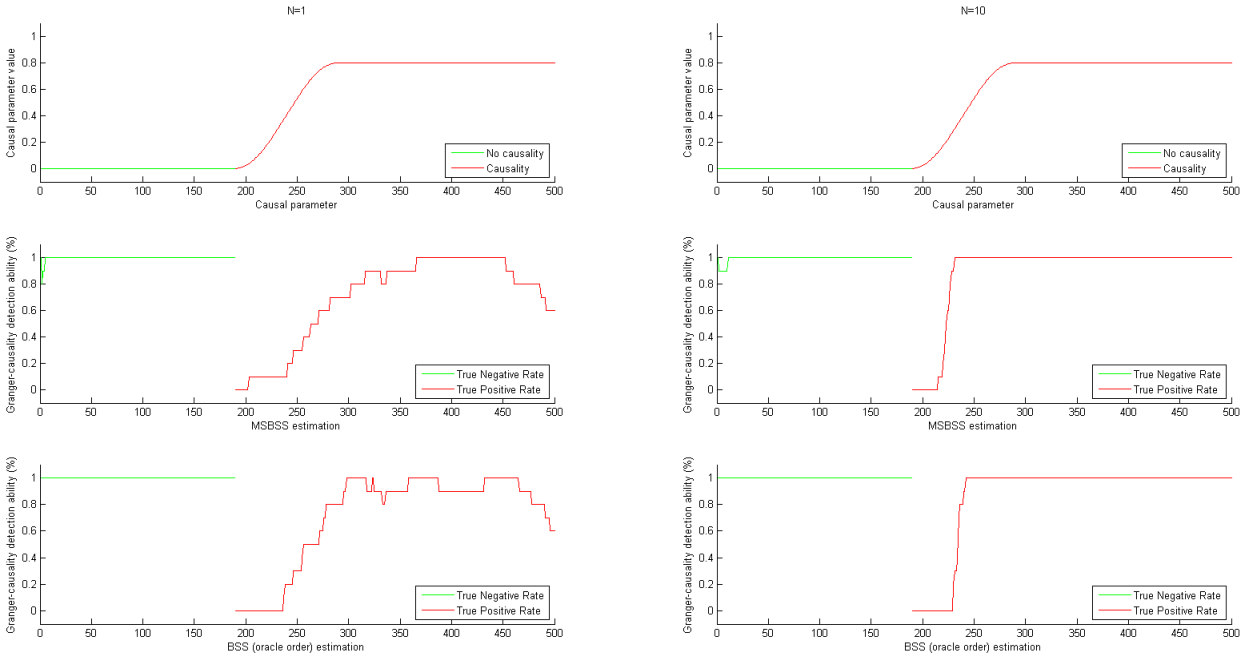

Figure M.34: Granger-causality detection ability for a non-normal error simulation of order 8 and series length 500 and causal parameter 0.8.

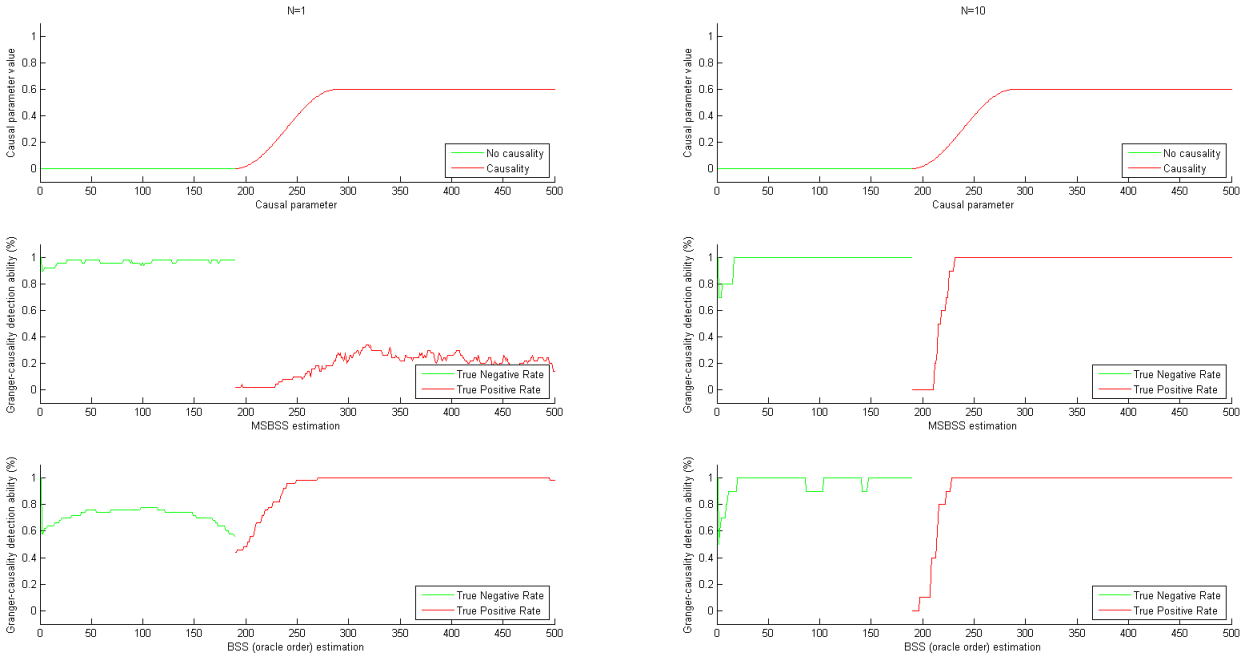

Figure M.35: Granger-causality detection ability for a non-normal error simulation of order 1 and series length 500 and causal parameter 0.6.

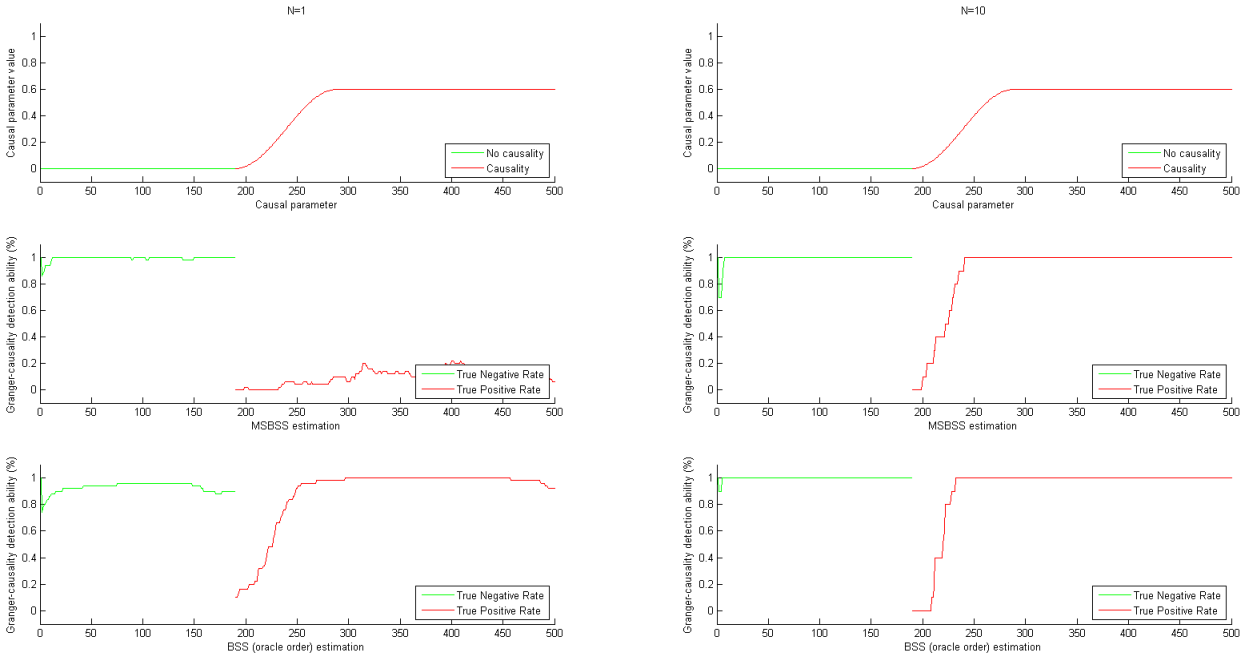

Figure M.36: Granger-causality detection ability for a non-normal error simulation of order 2 and series length 500 and causal parameter 0.6.

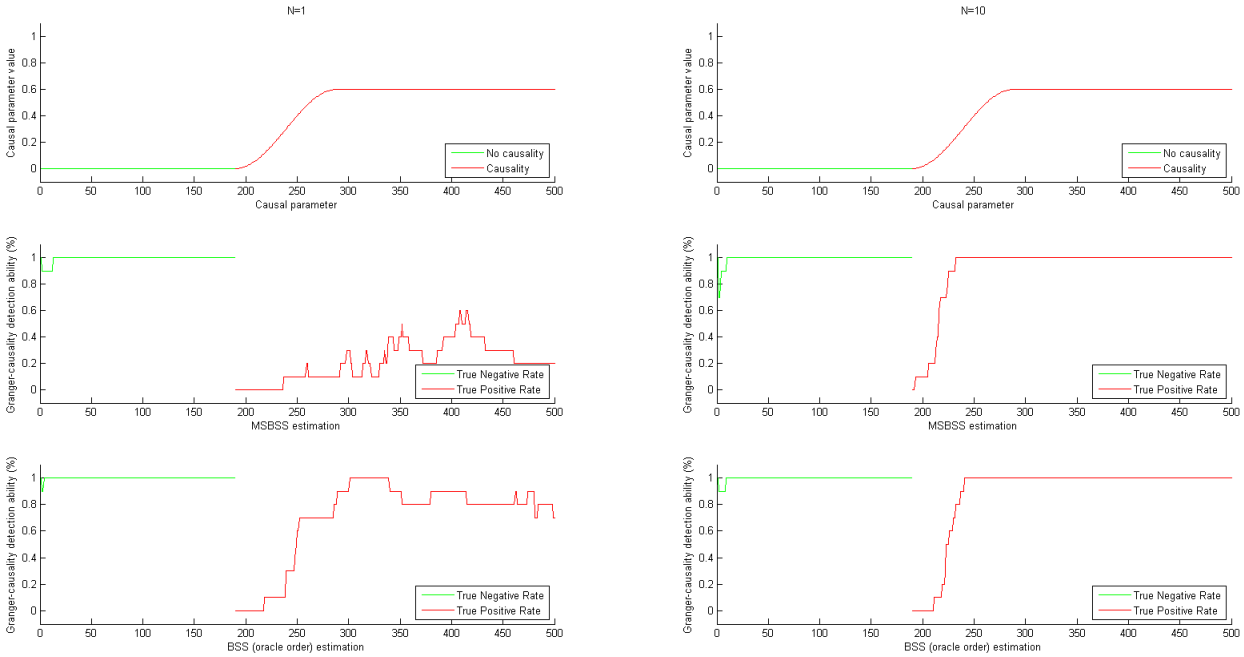

Figure M.37: Granger-causality detection ability for a non-normal error simulation of order 4 and series length 500 and causal parameter 0.6.

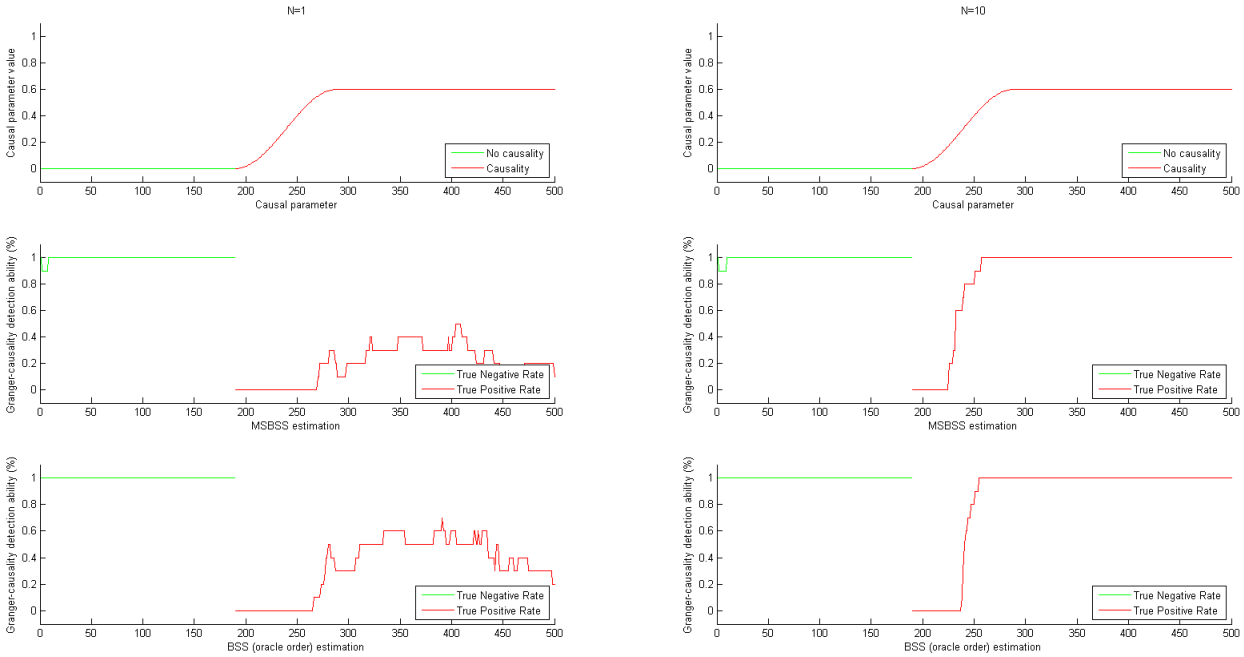

Figure M.38: Granger-causality detection ability for a non-normal error simulation of order 8 and series length 500 and causal parameter 0.6.

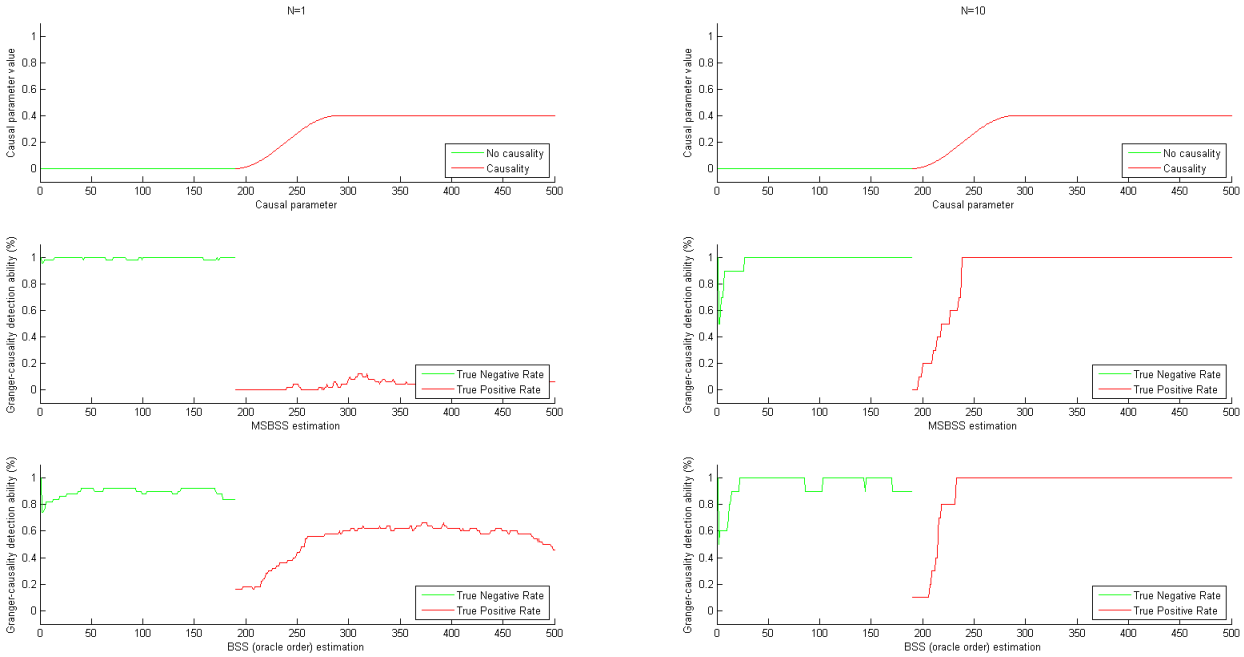

Figure M.39: Granger-causality detection ability for a non-normal error simulation of order 1, series length 500 and causal parameter 0.4. See Figure 2 for more details.

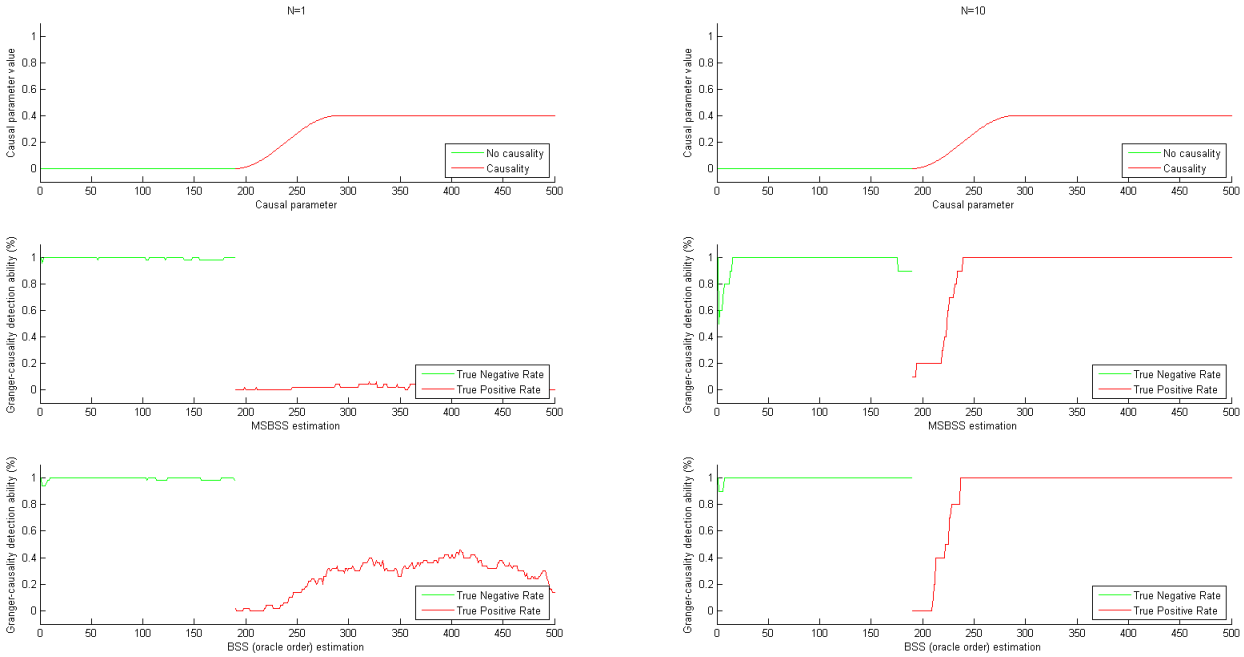

Figure M.40: Granger-causality detection ability for a non-normal error simulation of order 2 and series length 500 and causal parameter 0.4.

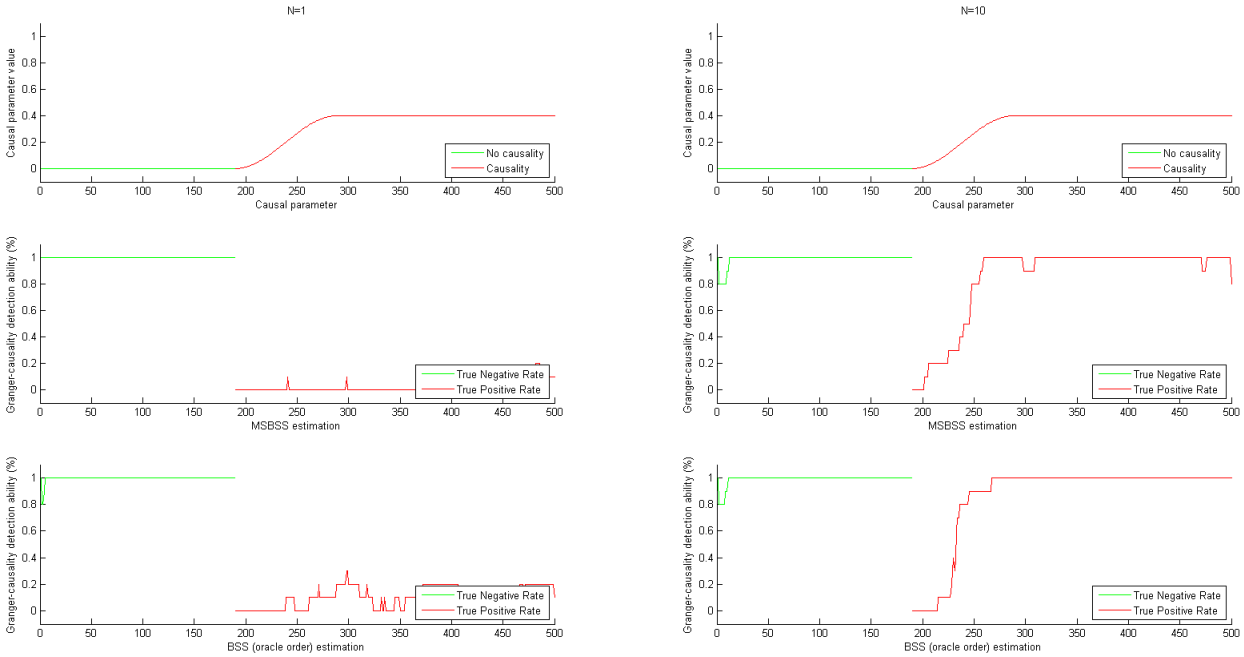

Figure M.41: Granger-causality detection ability for a non-normal error simulation of order 4 and series length 500 and causal parameter 0.4.

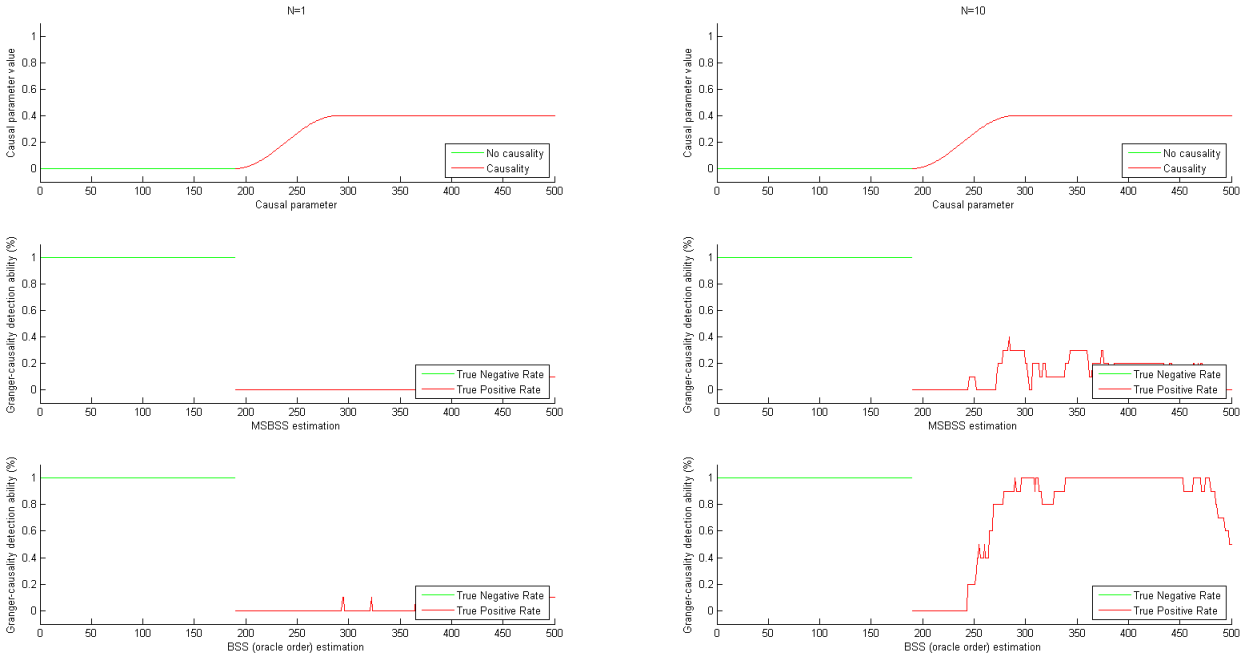

Figure M.42: Granger-causality detection ability for a non-normal error simulation of order 8 and series length 500 and causal parameter 0.4.

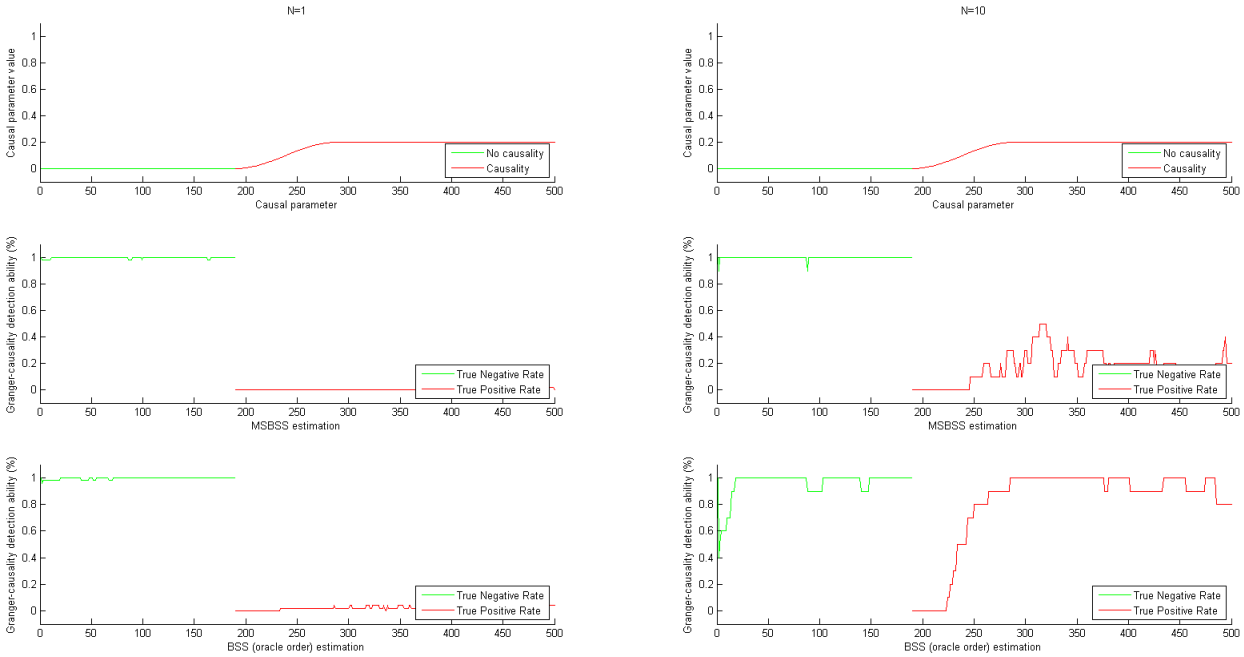

Figure M.43: Granger-causality detection ability for a non-normal error simulation of order 1 and series length 500 and causal parameter 0.2.

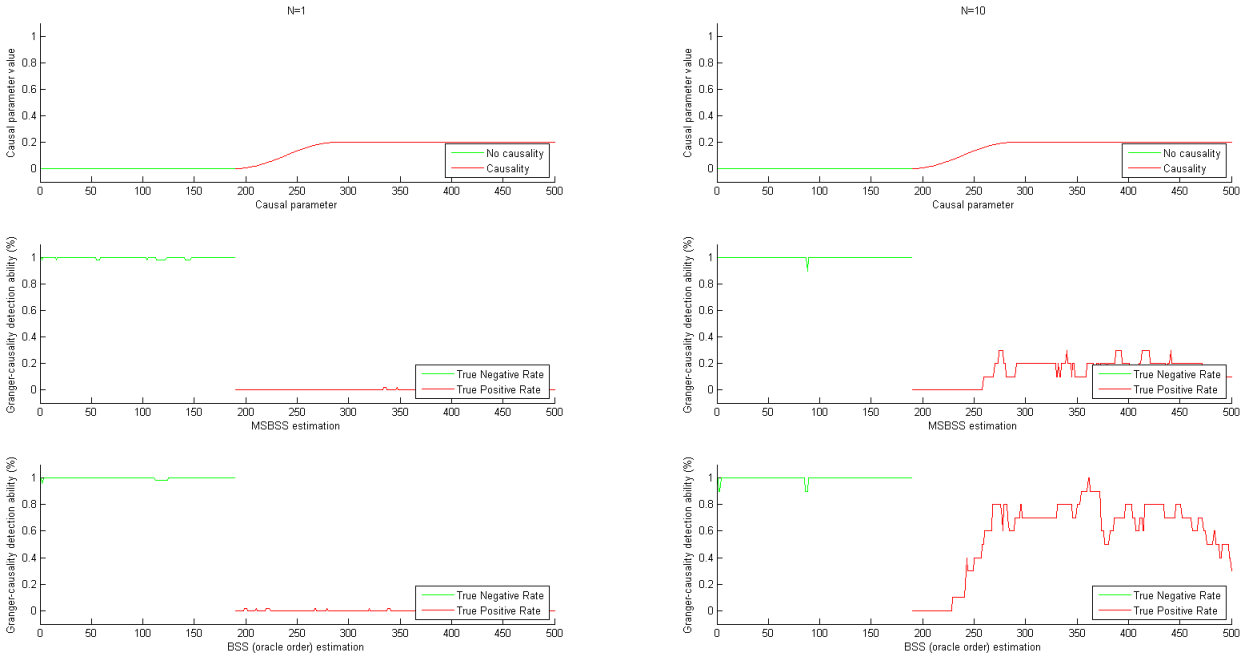

Figure M.44: Granger-causality detection ability for a non-normal error simulation of order 2 and series length 500 and causal parameter 0.2.

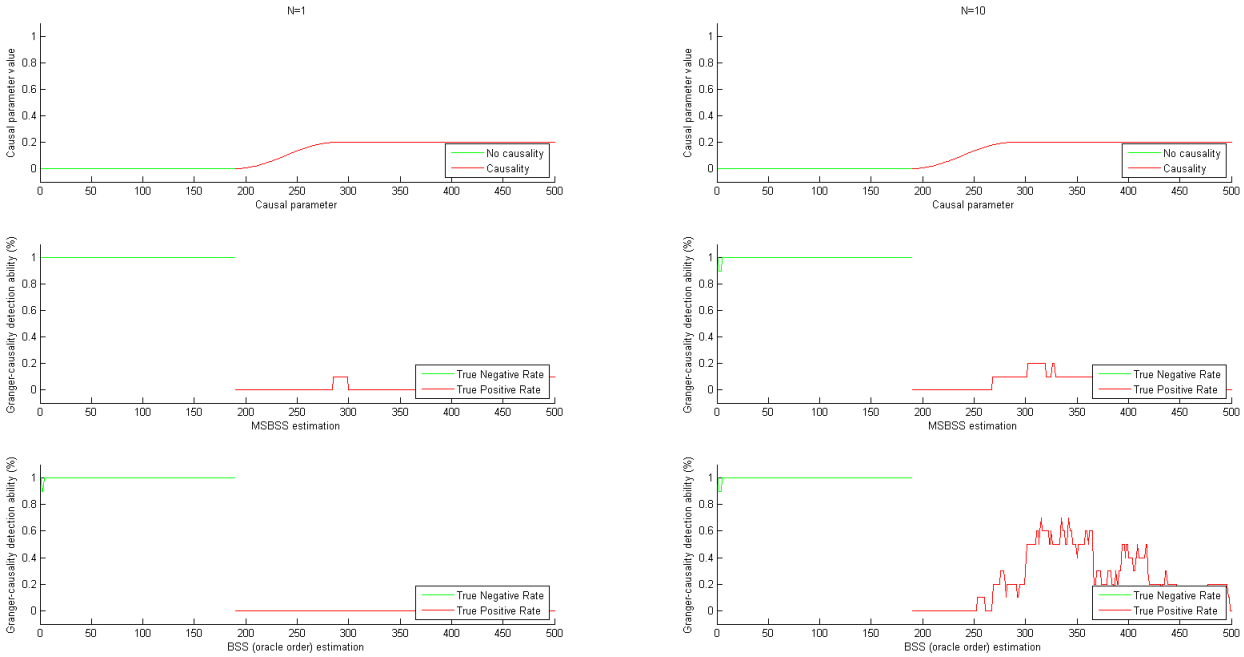

Figure M.45: Granger-causality detection ability for a non-normal error simulation of order 4 and series length 500 and causal parameter 0.2.

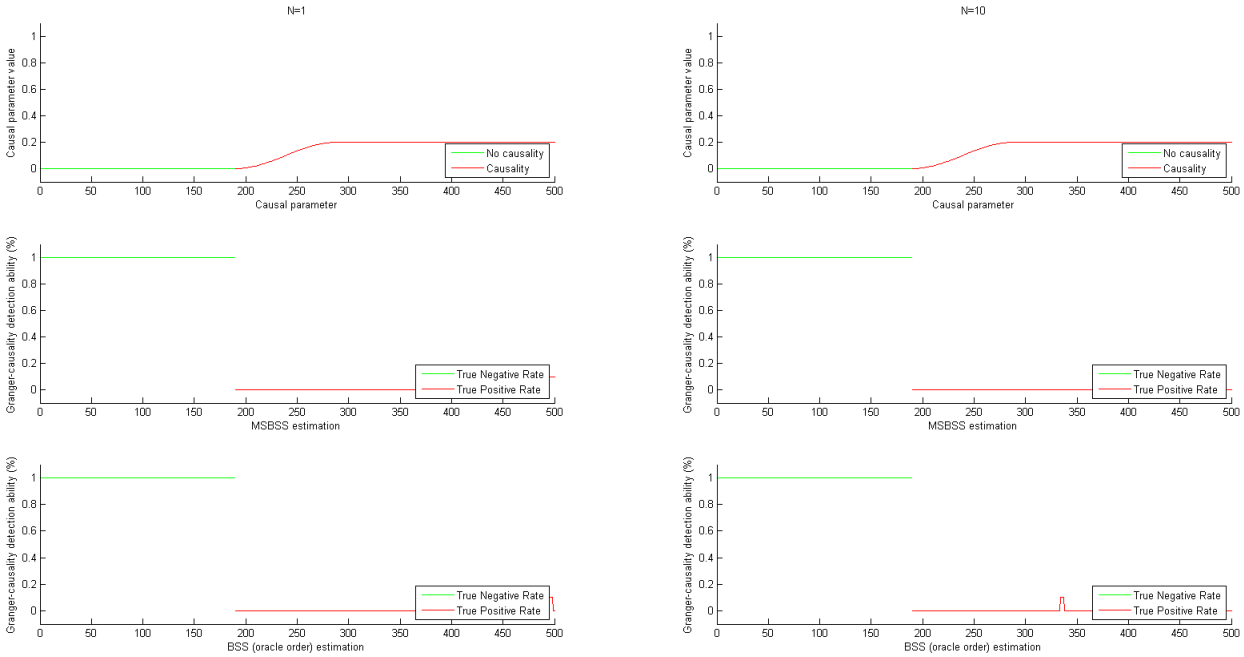

Figure M.46: Granger-causality detection ability for a non-normal error simulation of order 8 and series length 500 and causal parameter 0.2.